arXiv:1610.08637v1 [stat.ML] 27 Oct 2016

Statistical Inference for Model Parameters in Stochastic Gradient Descent

Xi Chen* Jason D. Lee† Xin T. Tong‡ and Yichen Zhang§

Abstract

The stochastic gradient descent (SGD) algorithm has been widely used in statistical estimation for large-scale data due to its computational and memory efficiency. While most existing work focuses on the convergence of the objective function or the error of the obtained solution, we investigate the problem of statistical inference of the true model parameters based on SGD. To this end, we propose two consistent estimators of the asymptotic covariance of the average iterate from SGD: (1) an intuitive plug-in estimator and (2) a computationally more efficient batch-means estimator, which only uses the iterates from SGD. As the SGD process forms a time-inhomogeneous Markov chain, our batch-means estimator with carefully chosen increasing batch sizes generalizes the classical batch-means estimator designed for time-homogenous Markov chains. The proposed batch-means estimator is of independent interest, which can be potentially used for estimating the covariance of other time-inhomogeneous Markov chains. Both proposed estimators allow us to construct asymptotically exact confidence intervals and hypothesis tests.

We further discuss an extension to conducting inference based on SGD for high-dimensional linear regression. Using a variant of the SGD algorithm, we construct a debiased estimator of each regression coefficient that is asymptotically normal. This gives a one-pass algorithm for computing both the sparse regression coefficient estimator and confidence intervals, which is computationally attractive and applicable to online data.

*Stern School of Business, New York University, Email: xichen@nyu.edu

†Marshall School of Business, University of Southern California, Email: jasondlee88@gmail.com

‡Department of Mathematics, National University of Singapore, Email: mattxin@nus.edu.sg

§Stern School of Business, New York University, Email: yzhang@stern.nyu.edu

1 Introduction

Estimation of model parameters by minimizing an objective function is a fundamental idea in statistics. Let $x^* \in \mathbb{R}^d$ be the true d-dimensional model parameters[1]. In many common models, x^* is the minimizer of a convex objective function $F(x)$ from $\mathbb{R}^d$ to $\mathbb{R}$ defined as follows,

$$F(x) := \mathbb{E}_{\zeta\sim\Pi} f(x,\zeta) = \int f(x,\zeta)\mathrm{d}\Pi(\zeta), \tag{1}$$

where ζ denotes the random sample from the probability distribution Π and $f(x,\zeta)$ is the loss function. In Section 2.2, we provide two classical estimation problems (i.e., linear regression and logistic regression) as motivating examples.

One fundamental problem of stochastic optimization is to compute the minimizer of the population function,

$$\min_x F(x) = \mathbb{E}_{\zeta\sim\Pi} f(x,\zeta). \tag{2}$$

A widely used optimization method for minimizing $F(x)$ is the *stochastic gradient descent (SGD)*, which has a long history in optimization (see, e.g., Robbins and Monro [1951], Polyak and Juditsky [1992], Nemirovski et al. [2009]). In particular, let x_0 denote the starting point. SGD is an iterative algorithm, where the n-th iterate x_n takes the following form,

$$x_n = x_{n-1} - \eta_n \nabla f(x_{n-1},\zeta_n). \tag{3}$$

The step size η_n is a decreasing sequence in n, ζ_n is the n-th sample randomly drawn from the distribution Π, and $\nabla f(x_{n-1},\zeta_n)$ denotes the gradient of $f(x,\zeta_n)$ with respect to x at $x = x_{n-1}$. The algorithm outputs either the last iterate x_n or the average iterate

$$\bar{x}_n = \frac{1}{n}\sum_{i=1}^{n} x_i \tag{4}$$

as the solution to the optimization problem in (2). When $\bar{x}_n$ is adopted as the solution, the algorithm is referred to as averaged SGD (ASGD) and such an averaging step is known as the Polyak-Ruppert averaging [Ruppert, 1988, Polyak and Juditsky, 1992]. SGD has many computational and storage advantages over traditional deterministic optimization methods. For example, SGD only uses one pass over the data and the per-iteration time complexity of SGD is $O(d)$, which is independent of the sample size. There is no need for SGD to store the dataset, and thus SGD naturally fits in the online setting, where

[1] We follow the standard notation in optimization literature, which usually denotes the global optimum of an optimization problem (corresponding to the true model parameters) by x^*.

each sample arrives sequentially (e.g., search queries or transactional data). In addition, ASGD is known to achieve the optimal convergence rate in $\mathbb{E}(F(\bar{x}_n) - F(x^*))$ when $F(x)$ is strongly convex with the rate of $O(1/n)$ [Rakhlin et al., 2012][2]. It has become the prevailing optimization method for many machine learning tasks [Srebro and Tewari, 2010], e.g., training deep neural networks.

Based on the simple SGD template in (3), there are a large number of variants developed in the optimization and statistical learning literature. Most existing works only focus on the convergence in terms of the objective function or the distance between the obtained solution and the true minimizer x^* of (2). However, the statistical inference of x^* based on SGD has largely remained unexplored. In this paper, we propose computationally efficient methods to conduct the statistical inference of x^*. The proposed method are based on a classical result of ASGD, which characterizes the limiting distribution of $\bar{x}_n$ in (4). In particular, let $A = \nabla^2 F(x^*)$ be the Hessian matrix of $F(x)$ at $x = x^*$ and S be the covariance matrix of $\nabla f(x^*, \zeta)$, i.e.,

$$S = \mathbb{E}\left([\nabla f(x^*, \zeta)][\nabla f(x^*, \zeta)]^T\right). \tag{5}$$

Note that $\mathbb{E}\nabla f(x^*, \zeta) = \nabla F(x^*) = 0$. Ruppert [1988] and Polyak and Juditsky [1992] showed that when F is strongly convex with a Lipschitz gradient, by choosing appropriate step sizes, $\sqrt{n}(\bar{x}_n - x^*)$ converges in distribution to a multivariate normal random vector, i.e.,

$$\sqrt{n}(\bar{x}_n - x^*) \Rightarrow \mathcal{N}(0, A^{-1} S A^{-1}). \tag{6}$$

This result shows that the solution $\bar{x}_n$ from ASGD is asymptotically efficient and attains the Cramér-Rao lower bound. To construct an asymptotically valid confidence interval (or equivalently, an asymptotically valid hypothesis test that controls the type I error at the nominal level), we only need to obtain a consistent estimator of the asymptotic covariance of $\sqrt{n}\bar{x}_n$, i.e., $A^{-1} S A^{-1}$.

We propose two approaches to estimate the asymptotic covariance of $\sqrt{n}\bar{x}_n$ based on SGD. An intuitive approach is the *plug-in* estimator. In particular, we propose a thresholding estimator $\widetilde{A}_n$ of A based on the sample estimate $A_n = \frac{1}{n}\sum_{i=1}^{n} \nabla^2 f(x_{i-1}, \zeta_i)$. Together with the sample estimate S_n of S, the asymptotic covariance $A^{-1} S A^{-1}$ is estimated by $\widetilde{A}_n^{-1} S_n \widetilde{A}_n^{-1}$, which is proven to be a consistent estimator. However, there are several drawbacks of the plug-in estimator:

[2] When $F(x)$ is non-smooth, Rakhlin et al. [2012] proposed a slight modification of the averaging step.

1. The plug-in estimator requires the computation of the Hessian matrix of the loss function f, which could be computationally expensive for some nonlinear models. Further, it requires the inverse of the matrix $\widetilde{A}_n$, which has a large computational cost of $O(d^3)$.

2. The plug-in estimator requires prior knowledge of a lower bound on the minimum eigenvalue of A to construct the threshold estimator $\widetilde{A}_n$. Further, to establish the consistency result, it requires an additional Lipschitz condition over the Hessian matrix of the loss function (see Assumption 3.1).

3. To form confidence intervals for coordinates of x^*, we are only interested in estimating diagonal elements of $A^{-1}SA^{-1}$. However, the plug-in estimator directly computes the entire estimator $\widetilde{A}_n^{-1}S_n\widetilde{A}_n^{-1}$, which is computationally wasteful. Furthermore when d is large, storing $\widetilde{A}_n$ and S requires more than d^2 bits, which is prohibitive and more than necessary since the output only requires the diagonals.

Given these drawbacks of the intuitive plug-in estimator, a natural question arises: can we estimate the asymptotic covariance *only using the iterates from SGD without requiring any additional information* (e.g., Hessian matrix of f)? We provide an affirmative answer to this question by proposing a computationally and memory efficient *batch-means* estimator. Basically, we split the sequence of iterates $\{x_1, x_2, \ldots, x_n\}$ into $M+1$ batches with the sizes $n_0, n_1, \ldots, n_M$:

$$\underbrace{\{x_{s_0}, \ldots, x_{e_0}\}}_{\text{0-th batch}}, \underbrace{\{x_{s_1}, \ldots x_{e_1}\}}_{\text{1-st batch}}, \ \ldots, \ \underbrace{\{x_{s_M}, \ldots, x_{e_M}\}}_{M\text{-th batch}}.$$

Here s_k and e_k are the starting and ending point of the k-th batch with $s_0 = 1$, $e_M = n$; and $n_k = e_k - s_k + 1$ is the batch size. The 0-th batch is discarded since the iterates in the 0-th batch are far from the optimum. For each batch k, we compute the mean of the iterates in the k-th batch and denote it by $\overline{X}_{n_k}$. The batch-means estimator is a "weighted" sample covariance matrix, which treats each batch mean as a sample:

$$\frac{1}{M}\sum_{k=1}^{M} n_k(\overline{X}_{n_k} - \overline{X}_M)(\overline{X}_{n_k} - \overline{X}_M)^T, \tag{7}$$

where $\overline{X}_M$ is the mean of all the iterates except for those from the 0-th batch. It is clear that the batch-means estimator does not require computing the Hessian matrix of f, nor any matrix inversion, thus is computationally more attractive than the plug-in estimator.

In addition, to form confidence intervals for coordinates of x^*, one only needs to compute and store the diagonal elements of batch-means estimator in (7).

The idea of batch-means estimator can be traced to Markov Chain Monte Carlo (MCMC), where the batch-means method with equal batch size (see, e.g., Glynn and Iglehart [1990], Damerdji [1991], Geyer [1992], Fishman [1996], Jones et al. [2006], Flegal and Jones [2010]) is widely used for variance estimation in a time-homogeneous Markov chain. The SGD iterates in (3) indeed form a Markov chain as x_n only depends on x_{n-1}. However, since the step size sequence η_n is a diminishing sequence, it is a *time-inhomogenous* Markov chain. Moreover, the asymptotic behavior of SGD and MCMC are fundamentally different, while the former converges to the optimum, the latter travels ergodically inside the state space. As a consequence of these two differences, the previous literature on batch-means methods are not applicable to our analysis. To address this challenge, our new batch-means method constructs batches of increasing size. The sizes of batches are chosen to ensure that the correlation decays appropriately among far apart batches, so that far apart batch means can be roughly treated as independent. In Theorem 3.2 below, the proposed batch-means method is shown to be a consistent estimator of the asymptotic covariance $A^{-1}SA^{-1}$. Further, we believe this new batch-means algorithm with increasing batch size is of independent interest since it can be used to estimate the covariance structure of other time-inhomogeneous Markov chains.

We further study the problem of confidence interval construction for x^* in high dimensional linear regression based on SGD, where the dimensionality d can be much larger than the sample size (or the number of iterates) n. It is noteworthy that the sample size n is also the number of iterates in SGD since each iteration of SGD samples a new data. In a high-dimensional setup, it is natural to solve an ℓ_1-regularized problem,

$$\min_x \; F(x) + \lambda \|x\|_1, \tag{8}$$

where $F(x)$ is defined in (1). A popular approach to solve (8) is the *proximal stochastic gradient approach* (see, e.g., Ghadimi and Lan [2012] and references therein) . However, due to the proximal operator (i.e., the soft-thresholding operator for ℓ_1-regularized problem), the distribution of the average iterate $\bar{x}_n$ will no longer converge to a multivariate normal distribution. To address this challenge, we use the recently proposed RADAR algorithm [Agarwal et al., 2012], which is a variant of SGD, together with the debiasing approach [Zhang and Zhang, 2014, van de Geer et al., 2014, Javanmard and Montanari, 2014]. The standard debiasing method relies on solving d convex optimization problems (e.g., node-wise Lasso in van de Geer et al. [2014]) to construct an approximation

of the inverse of covariance matrix of the design. Each optimization problem requires the computational cost of $O(nd^2T)$, where T is the number of iterations of the adopted deterministic optimization algorithm and generally in the hundreds to thousands. In contrast, we adopt the stochastic RADAR algorithm to solve these convex optimization problems, where each problem only requires one pass of the data with the computational cost $O(nd^2)$. Moreover, since the resulting approximate inverse covariance matrix from the stochastic RADAR is not an exact solution of the corresponding optimization problem, the analysis of van de Geer et al. [2014], which heavily relies on the KKT condition, is no longer applicable. We provide a new analysis to establish the asymptotic normality of the obtained estimator of x^* from the stochastic optimization algorithm.

1.1 Some related works on SGD

There is a large body of literature on stochastic gradient approaches and their applications to statistical learning problems (see, e.g., Zhang [2004], Nesterov and Vial [2008], Xiao [2010], Ghadimi and Lan [2012], Roux et al. [2012], Agarwal et al. [2012], Xiao and Zhang [2014] and references therein). Most works on SGD focus on the convergence rate of the objective function instead of the asymptotic distribution of the obtained solution. Thus, we only review a few closely related works with results on distributions.

Back in 1960's, Fabian [1968] studied the distribution of SGD iterates. However, without averaging, the asymptotic variance is inflated and thus the resulting statistical inference would have a reduced power even if the asymptotic is known. Ruppert [1988], Polyak and Juditsky [1992], Bach and Moulines [2011], Toulis and Airoldi [2016] studied the averaged SGD (ASGD) and its variant (e.g., implicit SGD in Toulis and Airoldi [2016]) and established the asymptotic normality and efficiency of the estimators. However these works do not discuss the estimation of the asymptotic covariance.

A few works in the SGD literature (e.g., Nesterov and Vial [2008], Nemirovski et al. [2009]) show large deviation results of $\Pr(\|\bar{x}_n - x^*\|_2 > t) \leq C(t)$ by combining the Markov inequality with the expected deviation of $\bar{x}_n$ to x^*. However, the ℓ_2-norm deviation bounds are not useful for constructing confidence intervals. Namely, the ℓ_2 error bounds are generally $O(\sigma\sqrt{\frac{d}{n}})$ (where σ^2 is the variance of the norm of the stochastic gradient), whereas a confidence interval for any single coordinate should be $O(\sigma\sqrt{\frac{1}{n}})$. Therefore, using an ℓ_2-norm error bound results in an extremely conservative confidence interval. Instead, we will use the central limit theorem that shows $\lim \Pr(|\bar{x}_{n,j} - x^*| > z_{q/2}\sigma/\sqrt{n}) \to 1 - q$, where $z_{q/2}$ is the $(1 - q/2)$-quantile of the standard normal distribu-

tion. This allows us to construct an asymptotically exact confidence interval that is also asymptotically efficient, in the sense that it attains the Cramér-Rao lower bound.

We also note that Toulis and Airoldi [2016] established the asymptotic normality for the averaged implicit SGD procedure, which has the same limiting distribution as ASGD. Therefore, as long as our key Lemma 2.6 also holds for averaged implicit SGD, our estimators (i.e., plug-in or batch-means) would also provide a consistent covariance estimate for averaged implicit SGD. We leave the verification of Lemma 2.6 for averaged implicit SGD as a future work direction.

1.2 Notations and organization of the paper

As a summary of notations, throughout the paper, we will use $\|x\|_p$ to denote the vector ℓ_p-norm of x, $\|X\|$ the matrix operator norm of X and $\|X\|_\infty$ the element-wise ℓ_∞-norm of X (i.e., $\|X\|_\infty = \max_{i,j} |X_{ij}|$). For a square matrix X, we denote its trace by $\mathrm{tr}(X)$. For a positive semi-definite (PSD) matrix A, let $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ be the maximum and minimum eigenvalue of A and when A is strictly positive definite, we use $\lambda_A > 0$ to denote any lower bound on $\lambda_{\min}(A)$. For a vector a of length d and any index subset $J \subseteq \{1, \ldots, d\}$, we denote by a_J the sub-vector of a with the elements indexed by J and a_{-J} the sub-vector of a with the elements indexed by $\{1, \ldots, d\} \backslash J$. Similarly for a $d_1 \times d_2$ matrix X and two index subsets $R \subseteq \{1, \ldots, d_1\}$ and $J \subseteq \{1, \ldots, d_2\}$, we denote by $X_{R,J}$ the $|R| \times |J|$ sub-matrix of X with elements in rows in R and columns in J. When $R = \{1, \ldots, d_1\}$ or $J = \{1, \ldots, d_2\}$, we denote $X_{R,J}$ by $X_{\cdot,J}$ or $X_{R,\cdot}$, respectively. We use I to denote the identity matrix. For any $a \in \mathbb{R}$, $(a)^+ = \max(a, 0)$. The function $\Phi(\cdot)$ denotes the CDF of the standard normal distribution.

Let $c, c_1, c_2, C, C_1, C_2, \ldots$, be the universal constants which could change from line to line. For any sequences $\{a_n\}$ and $\{b_n\}$ of positive numbers, we write $a_n \gtrsim b_n$ if $a_n \geq c b_n$ holds for all n and some absolute constant $c > 0$, $a_n \lesssim b_n$ if $b_n \gtrsim a_n$ holds, and $a_n \asymp b_n$ if both $a_n \gtrsim b_n$ and $a_n \lesssim b_n$ hold.

The rest of the paper is organized as follows. In Section 2, we discuss the problem setup and necessary assumptions, and also provide the convergence rate of iterates in Lemma 2.6. In Section 3, we propose the plug-in estimator and batch-means estimator for estimating the asymptotic covariance of $\bar{x}_n$ from ASGD. In Section 4, we discuss how to conduct inference for high-dimensional linear regression. In Section 5, we demonstrate the proposed methods by simulated experiments. All the proofs are given in the Appendix.

2 Problem Setup

Let us first introduce the problem setup and the assumptions made in this work.

2.1 Assumptions

In the classical work of Polyak and Juditsky [1992], the SGD method was introduced in a different formulation from (3): $x_n = x_{n-1} - \eta_n \nabla f(x_{n-1}, \zeta_n)$. In particular, the iteration was given by

$$x_n = x_{n-1} - \eta_n \nabla F(x_{n-1}) + \eta_n \xi_n. \tag{9}$$

This formulation decomposes the descent into two parts: $\nabla F(x_{n-1})$ represents the direction of population gradient which is the major driving force behind the convergence of SGD, and ξ_n is a sequence of martingale difference. In fact, the iteration in (9) is a more general formulation which includes (3) as a special case. To see this, we write (3) as

$$x_n = x_{n-1} - \eta_n \nabla F(x_{n-1}) + \eta_n \xi_n, \quad \xi_n := \nabla F(x_{n-1}) - \nabla f(x_{n-1}, \zeta_n), \tag{10}$$

where ξ_n is a martingale sequence due to Fubini's theorem:

$$\begin{aligned} \mathbb{E}_{n-1}[\xi_n] &= \nabla F(x_{n-1}) - \mathbb{E}_{n-1} \nabla f(x_{n-1}, \zeta_n) \\ &= \nabla F(x_{n-1}) - \nabla \mathbb{E}_{n-1} f(x_{n-1}, \zeta_n) = 0. \end{aligned}$$

Here and in the sequel, $\mathbb{E}_n(\cdot)$ denotes the conditional expectation $\mathbb{E}(\cdot|\mathcal{F}_n)$, where $\mathcal{F}_n$ is the sigma-algebra generated by $\{\zeta_1, \ldots, \zeta_n\}$ (note: ζ_k is the k-th sample). Let $\Delta_n := x_n - x^*$ be the error of the n-th iterate. It it noteworthy that by subtracting x^* from both sides of (10), the recursion (10) is equivalent to

$$\Delta_n = \Delta_{n-1} - \eta_n \nabla F(x_{n-1}) + \eta_n \xi_n, \tag{11}$$

which will be extensively used throughout the paper.

We first make the following standard assumptions on the population loss function F in (1).

Assumption 2.1 (Strong Convexity and Lipschitz continuity of the gradient)**.** *Assume that the objective function $F(x)$ is continuously differentiable and strongly convex with parameter $\mu > 0$, that is for any x_1 and x_2*

$$F(x_2) \geq F(x_1) + \langle \nabla F(x_1), x_2 - x_1 \rangle + \frac{\mu}{2} \|x_1 - x_2\|_2^2$$

Further, assume that ∇F is Lipschitz continuous with the constant L_F, i.e., for any x_1 and x_2,

$$\|\nabla F(x_1) - \nabla F(x_2)\|_2 \leq L_F \|x_1 - x_2\|_2. \tag{12}$$

We also assume F has a Hessian matrix A at $x = x^*$:

$$A = \nabla^2 F(x^*). \tag{13}$$

The strongly convexity of $F(x)$ guarantees the uniqueness of the minimizer x^*. Moreover, due to the strong convexity, we immediately find that $A \succeq \mu I$.

Let $\Delta_n := x_n - x^*$ be the error of the n-th iterate. We further assume that the martingale difference ξ_n satisfies the following conditions.

Assumption 2.2. *The followings hold for the sequence $\xi_n = \nabla F(x_{n-1}) - \nabla f(x_{n-1}, \zeta_n)$:*

1. *It is a martingale difference sequence: $\mathbb{E}_{n-1}\xi_n = 0$.*

2. *The conditional covariance of ξ_n has an expansion around $x = x^*$ of polynomial order $\kappa = 1$ or 2:*

$$\mathbb{E}_{n-1}\xi_n\xi_n^T = S + \Sigma(\Delta_{n-1}), \tag{14}$$

   and there exists a constant $\Sigma_2 > 0$ such that for any $\Delta \in \mathbb{R}^d$,

$$\|\Sigma(\Delta)\| \leq \Sigma_2(\|\Delta\|_2^\kappa + \|\Delta\|_2^2); \quad |\mathrm{tr}(\Sigma(\Delta))| \leq \Sigma_2(\|\Delta\|_2^\kappa + \|\Delta\|_2^2). \tag{15}$$

   Note that S is the covariance matrix of $\nabla f(x^*, \zeta)$ defined in (5).

3. *There exists a constant Σ_4 such that the fourth conditional moment of ξ_n is bounded by*

$$\mathbb{E}_{n-1}\|\xi_n\|_2^4 \leq \Sigma_4(1 + \|\Delta_{n-1}\|_2^4).$$

The constant κ in (15) plays an important role in the convergence rate of the covariance estimation. For example, when the loss is quadratic (e.g., the loss of linear regression, see Example 2.4 in Section 2.2), Eq. (15) holds for $\kappa = 2$, which produces a faster convergence rate in both theory and numerical simulation.

Assumption 2.2 is a mild condition over the regularity and boundedness of the loss function. In fact, one can easily verify Assumption 2.2 using the following lemma.

Lemma 2.3. *If there is a function $H(\zeta)$ with bounded fourth moment, such that the Hessian of $f(x,\zeta)$ is bounded by*

$$\left\|\nabla^2 f(x,\zeta)\right\| \le H(\zeta) \tag{16}$$

for all x, and $\nabla f(x^*,\zeta)$ have a bounded fourth moment, then Assumption 2.2 holds with $\kappa = 1$.

In the sequel, we will impose Assumptions 2.1 and 2.2. We note that under these two assumptions, the classical works by Ruppert [1988] and Polyak and Juditsky [1992] established the asymptotic normality and efficiency of the $\bar{x}_n$.

2.2 Two motivating examples

In this section, we illustrate Assumptions 2.1 and 2.2 on applications of SGD to two popular statistical estimation problems.

Example 2.4 (Linear Regression)**.** *Under the classical linear regression setup, let the n-th sample be $\zeta_n = (a_n, b_n)$, where the input $a_n \in \mathbb{R}^d$ is a sequence of random vectors independently drawn from the same multivariate distribution and the response $b_n \in \mathbb{R}$ follows a linear model,*

$$b_n = a_n^T x^* + \varepsilon_n. \tag{17}$$

*Here $x^* \in \mathbb{R}^d$ represents the true parameters of the linear model, and $\{\varepsilon_n\}$ are independently and identically distributed (*i.i.d.*) centered random variables, which are uncorrelated with a_n. For simplicity, we assume a_n and ε_n have all moments being finite. Given $\zeta_n = (a_n, b_n)$, loss function at x is a quadratic one:*

$$f(x,\zeta_n) = \frac{1}{2}(a_n^T x - b_n)^2.$$

and the true parameters $x^* = \operatorname{argmin}_x \left(F(x) := \mathbb{E} f(x,\zeta)\right)$. Given the loss function, the SGD iterates in (3) become,

$$x_n = x_{n-1} - \eta_n a_n (a_n^T x_{n-1} - b_n).$$

This can as well be written in the form of (9) as

$$x_n = x_{n-1} - \eta_n A \Delta_{n-1} + \eta_n \xi_n, \quad \xi_n := (A - a_n a_n^T)\Delta_{n-1} + a_n \varepsilon_n, \tag{18}$$

where $A = \mathbb{E} a_n a_n^T$ is the population gram matrix of a_n.

It is easy to find that

$$F(x) = \frac{1}{2}(x - x^*)^T A(x - x^*) + \mathbb{E}\varepsilon^2,$$

which implies that $\nabla F(x) = A(x - x^*)$ *and* $\nabla^2 F(x) = A$. *Therefore, Assumption 2.1 holds as long as* $\mu \le \lambda_{\min}(A) \le \lambda_{\max}(A) \le L_F$ *for some constants* $\mu, L_F > 0$. *This is commonly assumed in linear regression models.*

Lastly, we turn to the martingale difference sequence ξ_n. *First, we notice it is indeed a martingale sequence, since* $\mathbb{E}_{n-1}(A - a_n a_n^T) = 0$ *and* $\mathbb{E}_{n-1}\varepsilon_n = 0$. *Second, the conditional covariance takes the following form:*

$$\mathbb{E}_{n-1}\xi_n\xi_n^T = S + \Sigma(\Delta_{n-1}),$$

where

$$S := \mathbb{E}\varepsilon_n^2 a_n a_n^T, \quad \Sigma(\Delta_{n-1}) = \mathbb{E}_{n-1}[(a_n a_n^T - A)\Delta_{n-1}\Delta_{n-1}^T(a_n a_n^T - A)]. \tag{19}$$

Assumption 2.2 holds with $\kappa = 2$ *because*

$$\|\Sigma(\Delta)\| \le |\mathrm{tr}(\Sigma(\Delta))| \le \|\Delta\|_2^2 \mathbb{E}_{n-1}\left\|(a_n a_n^T - A)\right\|^2.$$

and by Hölder's inequality:

$$\begin{aligned} & \mathbb{E}_{n-1}\|(A - a_n a_n^T)\Delta_{n-1} + a_n\varepsilon_n\|_2^4 \\ \le\ & 8\|\Delta_{n-1}\|_2^4 \mathbb{E}_{n-1}\left\|A - a_n a_n^T\right\|^4 + 8\mathbb{E}_{n-1}(\varepsilon_n^4\|a_n\|_2^4). \end{aligned}$$

Example 2.5 (Logistic Regression)**.** *One of the most popular applications for general loss in statistics is the logistic regression for binary classification problems. In particular, the logistic model assumes that the binary response* $b_n \in \{-1, 1\}$ *is generated by the following probabilistic model,*

$$\Pr(b_n|a_n) = \frac{1}{1 + \exp\left(-b_n\langle a_n, x^*\rangle\right)},$$

while a_n *is an* i.i.d. *sequence generated by a fixed distribution. The population objective function is given by*

$$F(x) = \mathbb{E}f(x, \zeta_n) = \mathbb{E}\log(1 + \exp\left(-b_n\langle a_n, x\rangle\right)).$$

Let $\varphi(x) := \frac{1}{1+\exp(-x)}$ *denote the sigmoid function, we have*

$$\nabla F(x) = -\mathbb{E}\varphi(-b_n\langle a_n, x\rangle)b_n a_n,$$

Using the fact that $|\varphi'(x)| \le \frac{1}{4}$ *for any* x, $\nabla F(x)$ *is Lipschitz continuous with the Lipschitz constant* $L_F = \frac{1}{4}\mathbb{E}\|a_n\|_2^2$. *Moreover,*

$$A = \nabla^2 F(x^*) = \mathbb{E}\frac{a_n a_n^T}{[1+\exp(\langle a_n, x^*\rangle)][1+\exp(-\langle a_n, x^*\rangle)]}, \tag{20}$$

is clearly positive semi-definite. Additional non-degeneracy condition on a_n *(e.g., assuming* a_n *has a strictly positive density with respect to Lebesgue) will guarantee that* $\nabla^2 F(x)$ *has strictly positive minimum eigenvalue, see Lemma A.2 in Appendix.*

For the martingale difference assumption in Assumption 2.2, note that

$$\nabla^2 f(x, \zeta_n) = \frac{a_n a_n^T}{[1+\exp(\langle a_n, x\rangle)][1+\exp(-\langle a_n, x\rangle)]}, \tag{21}$$

which implies that $\left\|\nabla^2 f(x,\zeta_n)\right\| \le \|a_n\|_2^2$. *As long as* $\|a_n\|_2$ *has bounded eighth moment, Lemma 2.3 applies, which establishes Assumption 2.2.*

2.3 Inferences on SGD iterates

One of the fundamental theoretical results for SGD is that under Assumptions 2.1 and 2.2, the iterates x_n converge to the true minimizer x^*. Namely we have the following error moments bound for x_n:

Lemma 2.6. *Under Assumptions 2.1 and 2.2, the iterates of error* $\Delta_n = x_n - x^*$ *satisfy,*

1) There exist universal constants n_0 *and* C *such that for* $n > m \ge n_0$, *the followings hold:*

$$\mathbb{E}_m\|\Delta_n\|_2 \le \exp\left(-\tfrac{1}{4}\mu\sum_{i=m}^n \eta_i\right)\|\Delta_m\|_2 + Cm^{-\alpha/2},$$
$$\mathbb{E}_m\|\Delta_n\|_2^2 \le \exp\left(-\tfrac{1}{2}\mu\sum_{i=m}^n \eta_i\right)\|\Delta_m\|_2^2 + Cm^{-\alpha},$$
$$\mathbb{E}_m\|\Delta_n\|_2^4 \le \exp\left(-\tfrac{1}{2}\mu\sum_{i=m}^n \eta_i\right)\|\Delta_m\|_2^4 + Cm^{-2\alpha}.$$

2) Moreover, when the step size is chosen to be $\eta_n = \eta n^{-\alpha}$ *with* $\alpha \in (0,1)$, *the followings hold:*

$$\mathbb{E}\|\Delta_n\|_2 \lesssim n^{-\alpha/2}(1+\|\Delta_0\|_2),$$
$$\mathbb{E}\|\Delta_n\|_2^2 \lesssim n^{-\alpha}(1+\|\Delta_0\|_2^2),$$
$$\mathbb{E}\|\Delta_n\|_2^4 \lesssim n^{-2\alpha}(1+\|\Delta_0\|_2^4).$$

The proof of Lemma 2.6 is provided in Appendix. A similar result of in claim 1) of Lemma 2.6 has been shown in Theorem 3 in Bach and Moulines [2011] (see Theorem 3 and Appendix C). Here, we provide much simpler bounds on conditional moments of Δ_n. The claim 2) tells us how does the error decorrelates in terms of the number of iterations. By Lemma 2.6, we set the sequence of step sizes to be

$$\eta_n = \eta n^{-\alpha}, \qquad \alpha \in (0, 1), \tag{22}$$

throughout the paper.

The main objective of this paper is to construct a confidence interval of x^* based on the estimator $\overline{x}_n$ from ASGD. To this end, recall the Theorem 2 of Polyak and Juditsky [1992], which shows that

$$\sqrt{n} \cdot \overline{\Delta}_n \sim \mathcal{N}(0, A^{-1} S A^{-1}),$$

where $\overline{\Delta}_n = \frac{1}{n}\sum_{i=1}^n \Delta_i = \bar{x}_n - x^*$. It is worth noting that this result does not require that the model is well-specified. In a mis-specified case, the asymptotic distribution of $\bar{x}_n$ is centered at x^*, where x^* is the unique point such that $\nabla F(x^*) = 0$. Due to possible model mis-specification, the asymptotic covariance $A^{-1}SA^{-1}$ is of the sandwich covariance form (e.g., see Buja et al. [2013]); however, the sandwich covariance simplifies to the fisher information matrix if the model is indeed well-specified.

In order to constract a confidence interval for each coordinate of x^*, it suffices to estimate the asymptotic covariance matrix $A^{-1}SA^{-1}$, and form the interval $\bar{x}_{n,j} \pm z_{q/2}\hat{\sigma}_{jj}$, where $\hat{\sigma}_{jj}$ is an estimator of $(A^{-1}SA^{-1})_{jj}$ and $z_{q/2}$ is the $(1-q/2)$-quantile of the standard normal distribution (i.e., $z_{q/2} = \Phi^{-1}(1-q/2)$ and $\Phi(\cdot)$ is the CDF of the standard normal distribution).

3 Estimators for Asymptotic Covariance

Following the inference procedures illustrated above, it is essential to estimate the asymptotic covariance matrix $A^{-1}SA^{-1}$. In this section, we will propose two estimators, the plug-in estimator and the batch-means estimator.

3.1 Plug-in estimator

The idea of the plug-in estimator is to separately estimate A and S by some $\widehat{A}$ and $\widehat{S}$ and use $\widehat{A}^{-1}\widehat{S}\widehat{A}^{-1}$ as an estimator of $A^{-1}SA^{-1}$. Since x_i converges to x^* by Lemma 2.6,

according to the definitions of A and S in (13) and (5), an intuitive way to construct $\widehat{A}$ and $\widehat{S}$ is to use the sample estimate

$$A_n := \frac{1}{n}\sum_{i=1}^{n} \nabla^2 f(x_{i-1}, \zeta_i), \qquad S_n := \frac{1}{n}\sum_{i=1}^{n} \nabla f(x_{i-1}, \zeta_i)\nabla f(x_{i-1}, \zeta_i)^T,$$

as long as the information of $\nabla^2 f(x_{i-1}, \zeta_i)$ is available. Note that according to (3), $\nabla f(x_{i-1}, \zeta_i)$ is the i-th stochastic gradient. Since we are interested in estimating A^{-1}, it is necessary to avoid the possible singularity of A_n from statistical randomness. Therefore, we propose to use thresholding estimator $\widetilde{A}_n$, which has a strictly positive minimum eigenvalue. In particular, let λ_A be any lower bound on $\lambda_{\min}(A)$ so that $A \succeq \lambda_A I$. Further, let $\Psi D_n \Psi^T$ be the eigenvalue decomposition of A_n, where D_n is a non-negative diagonal matrix. We construct the thresholding estimator $\widetilde{A}_n$:

$$\widetilde{A}_n = \Psi \widetilde{D}_n \Psi^T, \quad \left(\widetilde{D}_n\right)_{i,i} = \max\left\{\frac{\lambda_A}{2}, (D_n)_{i,i}\right\}. \tag{23}$$

The threshold level is set to $\frac{\lambda_A}{2}$ for the ease of theoretical analysis. By the construction, it is guaranteed that $\lambda_{\min}\left(\widetilde{A}_n\right) \geq \frac{1}{2}\lambda_A$, and $\left\|\widetilde{A}_n^{-1}\right\|$ is bounded from above by $2\lambda_A^{-1}$, while clearly $\widetilde{A}_n = A_n$ if $A_n \succeq \frac{1}{2}\lambda_A I_d$.

With the construction of S_n and $\widetilde{A}_n$ in place, we propose the *plug-in estimator* as $\widetilde{A}_n^{-1} S_n \widetilde{A}_n^{-1}$. Our goal is to establish the consistency of the plug-in estimator, i.e.,

$$\mathbb{E}\left\|\widetilde{A}_n^{-1} S_n \widetilde{A}_n^{-1} - A^{-1} S A^{-1}\right\| \longrightarrow 0 \quad \text{as } n \to \infty.$$

Since this estimator relies on the Hessian matrix of the loss function, we need an additional assumption to establish the consistency,

Assumption 3.1. *There are constants L_2 and L_4 such that for all x,*

$$\mathbb{E}\|\nabla^2 f(x, \zeta) - \nabla^2 f(x^*, \zeta)\| \leq L_2 \|x - x^*\|_2,$$
$$\|\mathbb{E}[\nabla^2 f(x^*, \zeta_i)][\nabla^2 f(x^*, \zeta_i)]^T - AA^T\| \leq L_4.$$

We note that it is easy to verify that Assumption 3.1 holds for the two motivating examples in Section 2.2, because for quadratic loss, the Hessian matrix at any x is A itself, and (21) gives the Hessian for the logistic loss, which is Lipschitz in x and also bounded.

With this additional assumption, we first establish the consistency of the sample estimate A_n and S_n in the following lemma (its proof is provided in Appendix).

Lemma 3.2. *Under Assumptions 2.1, 2.2 and 3.1, the followings hold*

$$\mathbb{E}\|A_n - A\| \lesssim n^{-\frac{\alpha}{2}}, \quad \mathbb{E}\|S_n - S\| \lesssim n^{-\frac{\alpha}{2}},$$

where α is given in (22).

Using Lemma 3.2 and a matrix perturbation inequality for the inverse of a matrix (see Lemma B.1 in Appendix), we obtain the consistency result of the proposed plug-in estimator $\widetilde{A}_n^{-1} S_n \widetilde{A}_n^{-1}$:

Theorem 3.1 (Error rate of the plug-in estimator)**.** *Under Assumptions 2.1, 2.2 and 3.1, the plug-in estimator initialized from any bounded x_0 converges to the asymptotic covariance matrix,*

$$\mathbb{E}\left\|\widetilde{A}_n^{-1} S_n \widetilde{A}_n^{-1} - A^{-1} S A^{-1}\right\| \lesssim n^{-\frac{\alpha}{2}},$$

where α is given in (22).

We also note that since the element-wise ℓ_∞-norm is bounded from above by the matrix operator norm, so as $n \to \infty$,

$$\mathbb{E} \max_{ij} \left| \left(\widetilde{A}_n^{-1} S_n \widetilde{A}_n^{-1} - A^{-1} S A^{-1} \right)_{ij} \right| \to 0;$$

and thus $\left(A^{-1} S A^{-1}\right)_{jj}$ can be estimated by $\left(\widetilde{A}_n^{-1} S_n \widetilde{A}_n^{-1}\right)_{jj}$ for the construction of confidence intervals.

3.2 Batch-means estimator

Although the plug-in estimator is intuitive, as we discussed in the introduction, there are several drawbacks from both computational and theoretical perspectives, such as requiring the computation of the Hessian matrix, the inverse of $\tilde{A}_n$, the prior knowledge of $\lambda_{\min}(A)$ and an additional Assumption 3.1.

In this section, we develop the batch-means estimator, which only uses the iterates from SGD without requiring computation of any additional quantity. Intuitively, if all iterates are independent and share the same distribution, the asymptotic covariance can be directly estimated by the sample covariance,

$$\frac{1}{n} \sum_{i=1}^{n} (x_i - \bar{x})(x_i - \bar{x})^T.$$

Unfortunately, the SGD iterates are far from independent. To understand the correlation between two consecutive iterates, we note that for sufficiently large n such that x_{n-1} is close to x^*, by the Taylor expansion of $\nabla F(x_{n-1})$ at x^*, we have $\nabla F(x_{n-1}) \approx \nabla^2 F(x^*)(x_{n-1} - x^*) = A\Delta_{n-1}$. Combining this with the recursion in (11), we have for sufficiently large n,

$$\Delta_n \approx (I_d - \eta_n A)\Delta_{n-1} + \eta_n \xi_n. \tag{24}$$

Based on (24), the strength of correlation between Δ_n and Δ_{n-1} can be estimated by $\|I_d - \eta_n A\|$, which is very close to 1 as $\eta_n \asymp n^{-\alpha}$. To address the challenge of strong correlation among neighboring iterates, we split the entire sequence of iterates into batches with carefully chosen sizes and construct the batch-means estimator.

Batch-means estimators have been widely used in estimating the variance in time-homogeneous Markov chains and MCMC simulations (e.g., Damerdji [1991], Glynn and Iglehart [1990], Geyer [1992], Fishman [1996], Flegal and Jones [2010], Jones et al. [2006]). However, the traditional batch means estimator, see analysis in Jones et al. [2006], fails for the SGD process because the Markov chain formed by SGD iterates is not geometrically ergodic, and has a degenerate point mass for the invariant distribution. Moreover, due to the diminishing step sizes, the SGD iterates in (3) form a time-inhomogeneous Markov chain. We extend the classical batch-means estimator to the case of time-inhomogeneous Markov chains by using increasing batch sizes to ensure that the batches are sufficiently decorrelated. In particular, we split n iterates of SGD $\{x_1, \ldots, x_n\}$ into $M+1$ batches with the sizes $n_0, n_1, \ldots, n_M$:

$$\underbrace{\{x_{s_0}, \ldots, x_{e_0}\}}_{\text{0-th batch}}, \underbrace{\{x_{s_1}, \ldots x_{e_1}\}}_{\text{1-st batch}}, \ \ldots, \ \underbrace{\{x_{s_M}, \ldots, x_{e_M}\}}_{M\text{-th batch}}.$$

Here s_k and e_k are the starting and ending index of k-th batch with $s_0 = 1$, $s_k = e_{k-1} + 1$, $n_k = e_k - s_k + 1$, and $e_M = n$. We treat the 0-th batch as the "burn-in stage". More precisely, the iterates $\{x_{s_0}, \ldots, x_{e_0}\}$ will not be used for constructing the batch-means estimator since the step sizes are not small enough and the corresponding iterates in the 0-th batch are far away from the optimum. Let us denote the mean of the iterates for the k-th batch by $\overline{X}_{n_k}$ and the mean of all the iterates expect for $\{x_{s_0}, \ldots, x_{e_0}\}$ by $\overline{X}_M$, i.e.,

$$\overline{X}_{n_k} := \frac{1}{n_k} \sum_{i=s_k}^{e_k} x_i, \qquad \overline{X}_M := \frac{1}{e_M - e_0} \sum_{i=s_1}^{e_M} x_i. \tag{25}$$

The batch-means estimator is given by the following:

$$\frac{1}{M} \sum_{k=1}^{M} n_k (\overline{X}_{n_k} - \overline{X}_M)(\overline{X}_{n_k} - \overline{X}_M)^T. \tag{26}$$

The batch-means estimator has several advantages over the plug-in estimator. It only uses the iterates from SGD and does not require computing $\nabla^2 f$ or any matrix inversion, and thus is computationally more attractive.

Intuitively, the reason why our batch-means estimator with increasing batch size can overcome strong dependence between iterates is as follows. Although the correlation

between neighboring iterates is strong, it decays exponentially fast for far apart iterates. Roughly speaking, from (24), for large j and k, the strength of correlation between Δ_j and Δ_k can be estimated by

$$\prod_{i=j}^{k-1} \|I_d - \eta_{i+1}A\| \approx \exp\left(-\lambda_{\min}(A)\sum_{i=j}^{k-1}\eta_{i+1}\right). \tag{27}$$

Therefore, the correlations between the batch means $\overline{X}_{n_k}$ are close to zero if the batch sizes are large enough, in which case different batch means can be roughly treated as independent. As a consequence, the sample covariance gathered from the batch means will serve as a good estimator of the true asymptotic covariance.

The remaining difficulty is how to determine the batch sizes. The approximation of correlation given by (27) provides us a clear clue. If we want the correlation between two neighboring batches to be on the order of $\exp(-cN)$, where N is a parameter controlling the amount of decorrelation and c is a constant, we need

$$\sum_{i=s_k}^{e_k} \eta_i \asymp N$$

to hold for every batch k. When $\eta_i = \eta i^{-\alpha}$, $\sum_{i=s_k}^{e_k}\eta_i \asymp (e_k^{1-\alpha} - e_{k-1}^{1-\alpha})$, which leads to the following batch size setting:

$$e_k = ((k+1)N)^{\frac{1}{1-\alpha}} \quad k = 0, \ldots, M. \tag{28}$$

Given the total number of iterates n and noting that $e_M = n$, the decorrelation strength factor N takes the following form,

$$N = \frac{n^{1-\alpha}}{M+1}, \tag{29}$$

where M is the number of batches.

Under this setting, the batch-means covariance estimator (26) is consistent as shown in the following theorem.

Theorem 3.2. *Under Assumptions 2.1 and 2.2, the batch-means estimator initialized at any bounded x_0 is a consistent estimator for $\frac{1}{2} \le \alpha < 1$. In particular, for sufficiently large N and M, we have*

$$\begin{aligned}&\mathbb{E}\left\|M^{-1}\sum_{k=1}^{M} n_k(\overline{X}_{n_k} - \overline{X}_M)(\overline{X}_{n_k} - \overline{X}_M)^T - A^{-1}SA^{-1}\right\| \\ \lesssim\ & N^{-\frac{\kappa\alpha}{4-4\alpha}} M^{(\frac{1}{2}-\frac{\kappa\alpha}{4-4\alpha})^+ - \frac{1}{2}} + M^{-\frac{1}{2}} + N^{-\frac{1}{2}} + M^{(\frac{1-2\alpha}{2-2\alpha})^+ - \frac{1}{2}} N^{\frac{1-2\alpha}{2-2\alpha}},\end{aligned} \tag{30}$$

where the constant κ ($\kappa = 1$ or 2) is defined in Assumption 2.2.

As the number of iterates $n \to \infty$, by (29), we can choose $M, N \to \infty$ and thus the right hand side of (30) will converge to zero for any $\alpha \in [1/2, 1)$. The number of batches M can be possibly obtained by minimizing the right hand side of (30) numerically. However, in practice, this might not lead to satisfactory empirical performance. As we will show in the simulations in Section 5, wide choices between $M = n^{0.2}$ to $M = n^{0.3}$ lead to reasonably good and comparable coverage rates.

To motivate our proof, let us first consider the recursion (18) in the simple case of linear model with a quadratic loss function. By subtracting x^* from both sides of (18) and replacing Δ_n in (18) by a different notation U_n, we have

$$U_n = U_{n-1} - \eta_n A U_{n-1} + \eta_n \xi_n, \quad U_0 = \Delta_0. \tag{31}$$

The high-level idea of the proof consists of two steps:

1. establishing the consistency (and the rate of convergence) of the batch-means estimator based on the sequence U_n.

2. quantifying the difference between Δ_n and U_n, where Δ_n in (11) is the sequence of interest generated from SGD for a general loss function.

The details of this high-level proof is elaborated as follows. In fact, the so-called "oracle iterate sequence" in (31) has also been considered in Polyak and Juditsky [1992], which can be written in a more explicit form:

$$U_n = \prod_{k=1}^{n} (I - \eta_k A)\, U_0 + \sum_{m=1}^{n} \prod_{k=m+1}^{n} (I - \eta_k A)\, \eta_m \xi_m. \tag{32}$$

Given the sequence U_n, we construct the batch-means estimator based on U_n as follows

$$\frac{1}{M} \sum_{k=1}^{M} n_k (\overline{U}_{n_k} - \overline{U}_M)(\overline{U}_{n_k} - \overline{U}_M)^T, \tag{33}$$

where $\overline{U}_{n_k}$ and $\overline{U}_M$ are defined as in (25) with x_i being replaced by U_i. The analysis of the batch-means estimator from U_n is relatively simpler than that from x_n since the expression of U_n in (32) only involves the product of matrices and the martingale differences ξ_m. In particular, we establish the consistency of the batch-means estimator based on U_n in the following lemma (whose proof is provided in Appendix).

Lemma 3.3. *Under Assumptions 2.1 and 2.2, the batch-means estimator based on U_n with any U_0 satisfies the following inequality for sufficiently large N and M,*

$$\begin{aligned}&\mathbb{E}\left\|M^{-1}\sum_{k=1}^{M} n_k(\overline{U}_{n_k}-\overline{U}_M)(\overline{U}_{n_k}-\overline{U}_M)^T - A^{-1}SA^{-1}\right\| \\ \lesssim\ & N^{-\frac{\alpha}{4-4\alpha}}M^{(\frac{1}{2}-\frac{\kappa\alpha}{4-4\alpha})^+-\frac{1}{2}} + M^{-\frac{1}{2}} + N^{-\frac{1}{2}}.\end{aligned}$$

With Lemma 3.3 in place, to obtain the desired consistency result in Theorem 3.2, we only need to study the difference between Δ_n and U_n. In particular, let $\delta_n := \Delta_n - U_n$. We have the following recursion:

$$\begin{aligned}\delta_n &= \Delta_{n-1} - U_{n-1} - \eta_n \nabla F(x_{n-1}) - \eta_n A U_{n-1} \\ &= \delta_{n-1} - \eta_n A\delta_{n-1} + \eta_n(A\Delta_{n-1} - \nabla F(x_{n-1})).\end{aligned}$$

Notably, by replacing ξ_n in (31) with $A\Delta_{n-1} - \nabla F(\Delta_{n-1})$, δ_n follows a similar recursion relationship to that of the sequence U_n. Based on this observation, we show that δ_n is a sequence of small numbers, and hence Δ_n and U_n are close to each other. Combining this with Lemma 3.3, we will reach the conclusion in Theorem 3.2. The proof of Theorem 3.2 based on Lemma 3.3 is provided in Appendix.

4 High-dimensional Linear Regression

In Section 3.1 and 3.2, we assumed that the dimension d is fixed while $n \to \infty$. However in high-dimensional settings, it is often the case that $d \asymp n$ or $n = o(d)$. To conduct the inference for high-dimensional linear regression based on SGD, our approach combines the idea from the debiased estimator of Lasso [van de Geer et al., 2014, Javanmard and Montanari, 2014, Zhang and Zhang, 2014] and the recently proposed Regularization Annealed epoch Dual AveRaging (RADAR) algorithm [Agarwal et al., 2012].

Let us first briefly introduce some necessary background of the RADAR algorithm [Agarwal et al., 2012]. RADAR is a multi-epoch stochastic gradient based algorithm for solving the ℓ_1-regularized problem

$$\min_{x\in\mathbb{R}^d} F(x) + \lambda\|x\|_1.$$

Similar to the vanilla SGD, at each iteration t, the only required information of RADAR given the random data point ζ_t is the stochastic gradient $\nabla f(x_{t-1}, \zeta_t)$. However, to obtain the optimal convergence rate for a high-dimensional problem when the optimum x^* is

sparse, RADAR adopts a multi-epoch strategy. In particular, given the total number of n iterations, RADAR splits n iterations into K_n epochs, where each epoch has the length T_i for $i = 1, \ldots, K_n$ such that $\sum_{i=1}^{K_n} T_i = n$. For each epoch i, RADAR specifies (1) a regularization weight λ_i, which is different from λ, (2) a prox-center y_i with $y_1 = 0$, (3) a positive radius R_i such that $R_i = R_{i-1}/\sqrt{2}$ with R_1 being any upper bound on $\|x^*\|_1$. For each epoch i with the given prox-center y_i, RADAR runs T_i rounds of the stochastic dual averaging algorithm Nesterov [2009] (which is a variant of SGD) on the following regularized objective function:

$$\min_{x:\|x-y_i\|_p \le R_i} F(x) + \lambda_i \|x\|_1, \qquad p = \frac{2\log d}{2\log d - 1}, \tag{34}$$

and update the prox-center y_i to y_{i+1}. Agarwal et al. [2012] showed that the ℓ_1-distance between updated y_{i+1} and x^* and bounded from above by R_{i+1}, i.e., $\|y_{i+1} - x^*\|_1 \le R_{i+1}$. By setting the final solution $\widehat{x}_n = y_{K_n+1}$ and according to the geometric decaying of the sequence R_i,

$$\|\widehat{x}_n - x^*\|_1 \le \frac{R_1}{2^{K_n/2}},$$

which establishes the geometric convergence rate of $\widehat{x}_n$. Please refer to Agarwal et al. [2012] for more details of the RADAR algorithm.

In this paper, we only consider applying RADAR to high-dimensional linear model with the quadratic loss, i.e., $b = a^T x^* + \varepsilon$, where $x^* \in \mathbb{R}^d$ is sparse vector with $\|x^*\|_0 \le s$ and $\varepsilon \sim N(0, \sigma^2)$. For simplicity, we assume that σ is known. In practice, one can always replace it by a consistent estimator. Following Assumption (A2) in van de Geer et al. [2014], we also assume that the covariate a follows a d-dimensional multivariate normal distribution $a \sim N(0, A)$. By Assumption 2.1, we have both the maximum eigenvalue and minimum eigenvalue of A bounded above from infinity and below from zero, i.e.,

$$0 < \mu < \lambda_{\min}(A) < \lambda_{\max}(A) < L_F \tag{35}$$

Let us denote by $\Omega = A^{-1}$ the inverse covariance matrix and s_j the sparsity with respect to the j-th row of Ω, i.e.,

$$s_j = |\{1 \le k \le d : k \ne j, \Omega_{j,k} \ne 0\}|. \tag{36}$$

Our idea to conduct statistical inference of x^* is to utilize the debiased Lasso estimator [van de Geer et al., 2014, Javanmard and Montanari, 2014, Zhang and Zhang, 2014], which requires an estimate of Ω. All these works estimate Ω by solving deterministic convex optimization problems (e.g., node-wise Lasso in van de Geer et al. [2014]) and

thus, their approaches cannot be directly applied to the stochastic setting considered in this work. To address this issue, we compute the estimators for both x^* and Ω using the stochastic gradient based algorithm RADAR. In particular, we apply RADAR to the ℓ_1-regularized linear regression problem,

$$\operatorname*{argmin}_{x\in\mathbb{R}^d} \mathbb{E}(b - a^T x)^2 + \lambda\|x\|_1 \tag{37}$$

and let $\widehat{x}_n$ be the solution output from RADAR with n iterations. Due to the ℓ_1-regularization, $\widehat{x}_n$ is a biased estimator of x^* and thus a post-debiasing step is needed. van de Geer et al. [2014], Javanmard and Montanari [2014], Zhang and Zhang [2014] suggested the following debiased estimator

$$\widehat{x}^d = \widehat{x}_n + \frac{1}{n}\widehat{\Omega}D^T(b - D\widehat{x}_n), \tag{38}$$

where $D = [a_1, \ldots, a_n]^T \in \mathbb{R}^{n\times d}$ is the design matrix, $b = [b_1, \ldots, b_n]^T \in \mathbb{R}^{n\times 1}$ is the response vector, and $\widehat{\Omega}$ is an estimator for Ω.

To construct $\widehat{\Omega}$, we adopt the node-wise Lasso approach from Meinshausen and Bühlmann [2006], van de Geer et al. [2014] but use stochastic optimization instead of deterministic optimization. We again apply RADAR to the following optimization problem for each dimension $1 \le j \le d$,

$$\operatorname*{argmin}_{\gamma^j\in\mathbb{R}^{d-1}} \mathbb{E}\|a_j - a_{-j}\gamma^j\|_2^2 + \lambda_j\|\gamma^j\|_1, \tag{39}$$

where $a \sim N(0, A)$, $a_j \in \mathbb{R}$ denotes the j-th component of a and $a_{-j} \in \mathbb{R}^{d-1}$ denotes the sub-vector of a without the j-th component. Let $\widehat{\gamma}^j$ be the solution from RADAR with the components $\widehat{\gamma}^j = \{\widehat{\gamma}^j_k : k = 1, \ldots, d, k \neq j\}$. Further, we construct the estimator of $1/\Omega_{j,j}$ by

$$\widehat{\tau}_j = \frac{1}{n}(D_{\cdot,j} - D_{\cdot,-j}\widehat{\gamma}^j)^T D_{\cdot,j}, \tag{40}$$

where $D_{\cdot,j}$ is the j-th column of the design matrix D and $D_{\cdot,-j}$ is the design submatrix without the j-th column. Given $\widehat{\gamma}^j$ and $\widehat{\tau}_j$, let us denote the estimator of diagonal elements of Ω by $\widehat{T} := \mathrm{diag}(1/\widehat{\tau}_1, \ldots, 1/\widehat{\tau}_d)$ and let

$$\widehat{C} := \begin{pmatrix} 1 & -\widehat{\gamma}^1_2 & \ldots & -\widehat{\gamma}^1_d \\ -\widehat{\gamma}^2_1 & 1 & \ldots & -\widehat{\gamma}^2_d \\ \vdots & \vdots & \ddots & \vdots \\ -\widehat{\gamma}^d_1 & -\widehat{\gamma}^d_2 & \ldots & 1 \end{pmatrix}$$

The matrix Ω is estimated by,

$$\widehat{\Omega} = \widehat{T}\widehat{C}. \tag{41}$$

It is noteworthy that although the proposed $\widehat{\Omega}$ is in the same form as the estimator for A^{-1} in van de Geer et al. [2014], our $\widehat{\gamma}^j$ is different from the one in van de Geer et al. [2014]. More precisely, our $\widehat{\gamma}^j$ is the solution from a stochastic gradient based algorithm while van de Geer et al. [2014] obtained $\widehat{\gamma}^j$ from a deterministic optimization. Given all these ingredients in place, we present the stochastic gradient based construction of confidence interval for x_j^* for $j \in \{1, \ldots, d\}$ in Algorithm 1. The hypothesis test can also be performed once the estimator of the asymptotic variance of $\widehat{x}_j^d$ is available (see Theorem 4.2). We note that the proposed method is computationally more efficient than the methods based on deterministic optimization. It only requires one pass of the data with the computational cost $O(nd^2)$ and is applicable to online data (in contrast to hundreds to thousands passes of data with deterministic optimization algorithms used in existing methods).

The reason why we use RADAR instead of the vanilla SGD is because RADAR provides the optimal convergence rate in terms of the ℓ_1-norm. In particular, as a corollary of Proposition 1 in Agarwal et al. [2012], we obtain the following proposition.

Proposition 4.1. *Using the same settings of algorithm parameters as Proposition 1 in Agarwal et al. [2012], the $\widehat{x}_n$ and $\widehat{\gamma}^j$ from Algorithm 1 satisfy*

$$\|\widehat{x}_n - x^*\|_1 \lesssim s\sqrt{\frac{\log d}{n}}, \tag{42}$$

with high probability. Further, for each $j = \{1, \ldots, d\}$, we have

$$\|\widehat{\gamma}^j + \Omega_{j,j}^{-1}(\Omega_{j,-j})^T\|_1 \lesssim s_j\sqrt{\frac{\log d}{n}}, \tag{43}$$

hold with high probability.

The proof of Proposition 4.1 is provided in Appendix. Given Proposition 4.1, we state the inference result in the next theorem. We note that although the statement of the following theorem is similar to Theorem 2.2 in van de Geer et al. [2014], the proof is more technically involved. The main challenge is that the existing analysis is starting from the KKT condition of the deterministic optimization for estimating Ω. However, we estimate Ω using the stochastic gradient based RADAR and thus the corresponding KKT condition no longer holds. Please refer to the proof in Appendix for more details.

Algorithm 1 Stochastic Optimization Based Confidence Interval Construction for High-dimensional Sparse Linear Regression

Inputs:

Regularization parameter $\lambda \asymp \sqrt{\log d/n}$, and $\lambda_j \asymp \sqrt{\log d/n}$ for each dimensional j, the noise level σ, confidence level $1-\alpha$.

for $t = 1$ to n **do**

Randomly sample the data (a_t, b_t) and update the design $D \leftarrow [D^T, a_t]^T$ and response $b \leftarrow [b^T, b_t]^T$.

Update x_t by running one iteration of RADAR on the optimization problem (37) using the stochastic gradient $(a_t^T x_{t-1} - b_t)a_t$,

for $j = 1$ to d **do**

Update γ_t^j by running one iteration of RADAR on the optimization problem (39) using the stochastic gradient $(a_{t,-j}^T \gamma_{t-1}^j - a_{t,j})a_{t,-j}$.

end for

end for

Let $\widehat{x}_n = x_n$ and $\widehat{\gamma}^j = \gamma_n^j$ for $j \in \{1, \ldots, d\}$ be the final outputs.

Construct the debiased estimator according to (38) with $\widehat{\Omega}$ defined in (41).

Outputs:

The $(1-\alpha)$ confidence interval for each x_j^* :

$\widehat{x}_j^d \pm z_{\alpha/2}\sigma\sqrt{(\widehat{\Omega}\widehat{A}\widehat{\Omega})_{jj}/n}$, where $\widehat{A} = \frac{1}{n}D^T D$

Theorem 4.2. *Under Assumption 2.1 and a sparsity assumption $s = o(\sqrt{n}/\log d)$ and $\max_j s_j \leq Cs$ for some constant C. For suitable choices of $\lambda \asymp \sqrt{\log d/n}$ and $\lambda_j \asymp \sqrt{\log d/n}$, we have for all $j \in \{1, \ldots, d\}$,*

$$\frac{1}{\sigma\sqrt{(\widehat{\Omega}\widehat{A}\widehat{\Omega}^T)_{jj}}}\sqrt{n}(\widehat{x}_j^d - x_j^*) \overset{d}{\Rightarrow} N(0,1), \tag{44}$$

where $\widehat{x}^d$ is the debiased estimator defined in (38), $\widehat{\Omega}$ is defined in (41) and the sample covariance matrix $\widehat{A} = \frac{1}{n}D^T D$.

Theorem 4.2 shows that $\frac{1}{\sigma\sqrt{(\widehat{\Omega}\widehat{A}\widehat{\Omega}^T)_{jj}}}\sqrt{n}(\widehat{x}_j^d - x_j^*)$ converges in distribution to $N(0,1)$ for any j, which verifies the correctness of the asymptotic pointwise confidence interval in the output of Algorithm 1 for x_j^*. We also point out that since both RADAR and node-wise Lasso can handle generalized linear models under mild assumptions, our stochastic gradient based approach can also be extended to provide confidence intervals for parameters of generalized linear models. We leave this as a future work.

5 Numerical Simulations

In this section, we investigate the empirical performance of the plug-in estimator and batch-means estimator of the asymptotic covariance matrix.

For our simulation study, denoting by (a_n, b_n) the n-th sample and x^* the true parameter of the model, we consider linear and logistic models, where d-dimensional $\{a_n\}$ are fixed *i.i.d.* realizations from $\mathcal{N}(0, \Sigma)$. In both models, we specify three different covariance matrices:

- Identity: $\Sigma = I_d$;
- Toeplitz: $\Sigma_{i,j} = 0.5^{|i-j|}$;
- Equi Corr: $\Sigma_{i,j} = 0.2$ for all $i \neq j$, $\Sigma_{i,i} = 1$ for all i.

The noise ε_n in linear regression (see (17)) is set to *i.i.d.* $N(0, \sigma^2)$ with $\sigma = 1$. The parameter α in the step size is chosen to be 0.5. All the reported results are obtained by taking average of $N_{\text{sim}} = 1000$ independent runs. We consider the finite sample behavior of the plug-in estimator and the batch-means estimator for the inference of each individual regression coefficient x_j, $j \in \{1, 2, \ldots, d\}$.

5.1 Low dimensional cases

In each case, we consider the sample size $n = 10^5$ and $d = 5$, 20, 100. For each model, the corresponding parameter x^* is a d-dimensional vector linearly spaced between 0 and 1. For the plug-in estimator (plug-in in short), we use the true minimum eigenvalue of A as λ_A in (23). For linear regression, the matrix A is known since $A = \Sigma$; while for logistic regression, we estimate A by replacing the expectation in (20) by the sample average over a large number of (1 million) fresh samples a_n. For the batch-means estimator (BM in short), we consider three different choices of the number of batches: $M = n^{0.2}$, $M = n^{0.25}$, and $M = n^{0.3}$.

The performance of an estimator is measured by the average coverage rate of the confidence intervals and the average length of the intervals for each individual coefficient. Let us denote by CI_j the two-sided confidence interval for x_j^*. By setting the nominal coverage probability $1 - q$ to 95%, we report empirical versions of

- Average Coverage Rate (Cov Rate): $d^{-1} \sum_{j \in \{1,\ldots,d\}} \mathbb{P}[x_j^* \in \text{CI}_j]$;
- Average Length (Avg Len): $d^{-1} \sum_{j \in \{1,\ldots,d\}} \text{length}(\text{CI}_j)$.

Table 1: Linear Regression: the average coverage rate and length of confidence intervals, for the nominal coverage probability 95%. The second column (Plug-in) corresponds to the plug-in estimator. The columns (BM: n^c for c=0.2, 0.25, and 0.3) correspond to the batch-means estimator with n^c number of batches.

	d	Plug-in	BM: $n^{0.2}$	BM: $n^{0.25}$	BM: $n^{0.3}$	Oracle
Identity Σ						
Cov Rate (%)	5	95.68	90.76	92.72	93.96	
Avg Len ($\times 10^{-2}$)	5	1.52	1.43	1.47	1.48	1.24
Cov Rate (%)	20	95.01	91.00	92.62	93.64	
Avg Len ($\times 10^{-2}$)	20	1.45	1.40	1.42	1.42	1.24
Cov Rate (%)	100	94.97	90.61	92.18	92.85	
Avg Len ($\times 10^{-2}$)	100	1.43	1.37	1.38	1.37	1.24
Toeplitz Σ						
Cov Rate (%)	5	95.42	90.41	92.75	93.88	
Avg Len ($\times 10^{-2}$)	5	1.85	1.81	1.84	1.83	1.53
Cov Rate (%)	20	94.88	90.20	92.08	92.97	
Avg Len ($\times 10^{-2}$)	20	1.85	1.78	1.80	1.80	1.58
Cov Rate (%)	100	95.07	90.17	91.49	91.68	
Avg Len ($\times 10^{-2}$)	100	1.85	1.73	1.74	1.69	1.60
Equi Corr Σ						
Cov Rate (%)	5	94.79	91.07	91.99	92.49	
Avg Len ($\times 10^{-2}$)	5	1.61	1.56	1.58	1.59	1.31
Cov Rate (%)	20	95.05	90.62	92.34	93.48	
Avg Len ($\times 10^{-2}$)	20	1.59	1.53	1.55	1.56	1.36
Cov Rate (%)	100	95.00	90.38	91.80	92.48	
Avg Len ($\times 10^{-2}$)	100	1.60	1.53	1.53	1.52	1.38

For each setting, we also report the oracle length of the confidence interval with respect to the true covariance matrix $A^{-1}SA^{-1}$.

For linear regression, the asymptotic covariance is $A^{-1}SA^{-1} = \sigma^2\Sigma^{-1} = \Sigma^{-1}$ and the oracle interval length for each coordinate j will be $\frac{2z_{q/2}(\Sigma^{-1})_{jj}}{\sqrt{n}}$. Table 1 shows the empirical performance of the plug-in and BM under linear models with three different design covariance matrices.

From Table 1, both of the plug-in and BM achieve good performance. The plug-in gives

Table 2: Logistic Regression: the average coverage rate and length of confidence intervals, for the nominal coverage probability 95%. The second column (Plug-in) corresponds to the plug-in estimator. The columns (BM: n^c for c=0.2, 0.25, and 0.3) correspond to the batch-means estimator with n^c number of batches.

	d	Plug-in	BM: $n^{0.2}$	BM: $n^{0.25}$	BM: $n^{0.3}$	Oracle
Identity Σ						
Cov Rate (%)	5	95.05	89.56	90.04	89.26	
Avg Len ($\times 10^{-2}$)	5	2.50	2.29	2.25	2.11	2.51
Cov Rate (%)	20	95.04	89.38	90.08	89.07	
Avg Len ($\times 10^{-2}$)	20	2.53	2.32	2.27	2.13	2.51
Cov Rate (%)	100	94.73	89.72	90.79	90.37	
Avg Len ($\times 10^{-2}$)	100	2.70	2.43	2.36	2.16	2.51
Toeplitz Σ						
Cov Rate (%)	5	94.90	89.04	90.58	90.36	
Avg Len ($\times 10^{-2}$)	5	3.09	2.91	2.88	2.77	3.06
Cov Rate (%)	20	95.25	89.43	90.22	89.35	
Avg Len ($\times 10^{-2}$)	20	3.24	3.04	2.98	2.81	3.17
Cov Rate (%)	100	94.87	89.75	90.81	90.56	
Avg Len ($\times 10^{-2}$)	100	3.35	3.12	3.09	2.97	3.28
Equi Corr Σ						
Cov Rate (%)	5	94.76	88.56	89.18	88.18	
Avg Len ($\times 10^{-2}$)	5	2.63	2.41	2.35	2.21	2.61
Cov Rate (%)	20	94.76	89.48	90.92	90.85	
Avg Len ($\times 10^{-2}$)	20	2.83	2.65	2.63	2.54	2.74
Cov Rate (%)	100	94.89	89.43	90.13	89.19	
Avg Len ($\times 10^{-2}$)	100	3.00	2.90	2.85	2.68	2.78

better average coverage rate than BM: the average coverage rates in all different settings are nearly 95%. However, the average length of plug-in is usually larger than that of BM and the corresponding oracle interval length. On the other hand, BM achieves about 92% coverage rate when $M = n^{0.25}$ or $M = n^{0.3}$. Moreover, since BM only uses the iterates from SGD, it is computationally more efficient than plug-in, which requires the computation of the Hessian matrix and the inverse of the matrix $\tilde{A}_n$.

We further consider the logistic regression. To provide an oracle interval length based

on the true asymptotic covariance $A^{-1}SA^{-1} = A^{-1}$, we estimate A in (20) by its empirical version $\widehat{A}$ using one million fresh samples and the oracle interval length of each coordinate j is computed as $\frac{2z_{q/2}(\widehat{A}^{-1})_{jj}}{\sqrt{n}}$. We provide the result in Table 2 for different design covariance matrices. From Table 2, the plug-in still achieves nearly 95% average coverage rate. The BM achieves about 90% coverage rate and the average length is usually smaller than the oracle length. Moreover, as d becomes larger, the interval lengths for both estimators increase. Fianlly, the performance of BM is insensitive to the choice of the number of batches M: different M's lead to comparable average coverage rate.

In summary, our empirical results are consistent with our theoretical results in Theorem 3.1 and Theorem 3.2 and both the plug-in and BM provide asymptotically exact confidence intervals.

5.2 High dimensional cases

In a high dimensional setting, we consider the sample size $n = 100$, and the dimension $d = 500$. The active set $S_0 = \{1, 2, \ldots, s_0\}$, where the cardinality $s_0 = |S_0| = 3$ or 15. The non-zero regression coefficients $\{x_j^*\}_{j \in S_0}$ are from a fixed realization of s_0 *i.i.d.* uniform distribution $U[0, c]$ with $c = 25$.

First, we consider the average coverage rate and the average length of the intervals for individual coefficients corresponding to the variables in either S_0 or S_0^c where $S_0^c = \{1, \ldots, d\} \backslash S_0$. By setting the nominal coverage probability $1 - q$ to 95%, we report empirical versions of

- Average Coverage Rate (Cov Rate) on $S_0 = s_0^{-1} \sum_{j \in S_0} \mathbb{P}[x_j^* \in \mathrm{CI}_j]$
- Average Coverage Rate (Cov Rate) on $S_0^c = (d - s_0)^{-1} \sum_{j \in S_0^c} \mathbb{P}[0 \in \mathrm{CI}_j]$
- Average Length (Avg Len) on $S_0 = s_0^{-1} \sum_{j \in S_0} \mathrm{length}(\mathrm{CI}_j)$;
- Average Length (Avg Len) on $S_0^c = (d - s_0)^{-1} \sum_{j \in S_0^c} \mathrm{length}(\mathrm{CI}_j)$;

Our experimental setup follows directly from van de Geer et al. [2014], and we provide the oracle length of the confidence intervals for comparison. For linear regression, the asymptotic covariance is $A^{-1}SA^{-1} = \sigma^2\Sigma^{-1} = \Sigma^{-1}$ and the oracle interval length for each coordinate j will be $\frac{2z_{q/2}(\Sigma^{-1})_{jj}}{\sqrt{n}}$. We provide the result in Table 3 for different design covariance matrices.

Table 3: High dimensional linear regression, the average coverage rate and length of confidence intervals, for the nominal coverage probability 95%

Measure	Identity Σ	Toeplitz Σ	Equi Corr Σ
$S_0 = \{1, 2, 3\}$			
Cov Rate S_0(%)	92.07	91.42	90.37
Avg Len S_0	0.383	0.403	0.363
Cov Rate S_0^c(%)	90.44	90.21	89.78
Avg Len S_0^c	0.382	0.412	0.372
Oracle Len	0.392	0.506	0.438
$S_0 = \{1, 2, \ldots, 15\}$			
Cov Rate S_0 (%)	90.48	89.84	89.45
Avg Len S_0	0.379	0.430	0.384
Cov Rate S_0^c (%)	88.43	86.79	87.12
Avg Len S_0^c	0.360	0.425	0.383
Oracle Len	0.392	0.506	0.438

For high dimensional linear regression, Algorithm 1 achieves good performance, especially in sparse settings ($s_0 = 3$). From Table 3, the average coverage rate is about 90%. For less sparse problems ($s_0 = 15$), we still achieve about 88% average coverage rate for different design covariance matrices. Inferences on the active sets S_0 are slightly better than those on S_0^c. The average length on both sets is slightly smaller than the oracle length.

In summary, our empirical results are consistent with our theoretical results in Theorem 4.2.

6 Conclusions and Future Works

This paper presents two consistent estimators of the asymptotic variance of the average iterate from SGD, especially a computationally more efficient batch-means estimator, which only uses iterates from SGD. With the proposed estimators, we are able to construct asymptotically exact confidence intervals and hypothesis tests.

We further discuss statistical inference based on SGD for high-dimensional linear regression. An extension to generalized linear models is an interesting problem for future work.

The seminal work by Toulis and Airoldi [2016] develops the averaged implicit SGD procedure and provides the characterization of the limiting distribution. It would be interesting to establish the consistency of the batch-means estimator based on iterates from implicit SGD.

Acknowledgements

The authors would like to thank John Duchi for early discussions on Markov processes and the relationship to stochastic gradient.

References

A. Agarwal, S. Negahban, and M. J. Wainwright. Stochastic optimization and sparse statistical recovery: Optimal algorithms for high dimensions. In *Advances in Neural Information Processing Systems*, 2012.

F. Bach and E. Moulines. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *Advances in Neural Information Processing Systems*, 2011.

A. Buja, R. Berk, L. Brown, E. George, M. Traskin, K. Zhang, and L. Zhao. A conspiracy of random x and model violation against classical inference in linear regression. Technical report, Citeseer, 2013.

F. C. Chang. Inversion of a perturbed matrix. *Applied mathematics letters*, 19(2):169–173, 2006.

H. Damerdji. Strong consistency and other properties of the spectral variance estimator. *Management Science*, 37:1424–1440, 1991.

V. Fabian. On asymptotic normality in stochastic approximation. *Ann. Math. Statist*, pages 1327–1332, 1968.

G. Fishman. *Monte Carlo: Concepts, Algorithms and Applications.* Chapman and Hall, 1996.

J. Flegal and G. Jones. Batch means and spectral variance estimators in Markov chain Monte Carlo. *Annals of Statistics*, 38(2):1034–1070, 2010.

C. Geyer. Practical markov chain monte carlo. *Statistical Science*, 7(4):473–483, 1992.

S. Ghadimi and G. Lan. Optimal stochastic approximation algorithms for strongly convex stochastic composite optimization i: A generic algorithmic framework. *SIAM Journal of Optimization*, 22(4):1469–1492, 2012.

W. Glynn and D. Iglehart. Simulation analysis using standardized time series. *Mathematics of Operations Research*, 15:1–16, 1990.

A. Javanmard and A. Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *Journal of Machine Learning*, 15(1):2869–2909, 2014.

G. L. Jones, M. Haran, B. S. Caffo, and R. Neath. Fixed-width output analysis for markov chain monte carlo. *Journal of the American Statistical Association*, 101(476): 1537–1547, 2006.

P.-L. Loh, M. J. Wainwright, et al. High-dimensional regression with noisy and missing data: Provable guarantees with nonconvexity. *Annals of Statistics*, 40(3):1637–1664, 2012.

N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics*, (3):1436–1462, 06 2006.

A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

Y. Nesterov. Primal-dual subgradient methods for convex problems. *Math Programming*, 120(1):221–259, 2009.

Y. Nesterov and J.-P. Vial. Confidence level solutions for stochastic programming. *Automatica*, 44(6):1559–1568, 2008.

B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838–855, 1992.

A. Rakhlin, O. Shamir, and K. Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *Proceedings of the International Conference on Machine Learning*, 2012.

H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22(3):400–407, 1951.

N. L. Roux, M. Schmidt, and F. Bach. A stochastic gradient method with an exponential convergence rate for strongly-convex optimization with finite training sets. In *Advances in Neural Information Processing Systems*, 2012.

D. Ruppert. Efficient estimations from a slowly convergent robbins-monro process. Technical report, Cornell University Operations Research and Industrial Engineering, 1988.

N. Srebro and A. Tewari. Stochastic optimization for machine learning. Tutorial at International Conference on Machine Learning, 2010.

P. Toulis and E. M. Airoldi. Asymptotic and finite-sample properties of estimators based on stochastic gradients. *Annals of Statistics (to appear)*, 2016.

S. van de Geer, P. Bühlmann, Y. Ritov, and R. Dezeure. On asymptotically optimal confidence regions and tests for high-dimensional models. *Annals of Statistics*, 42(3): 1166–1202, 2014.

L. Xiao. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning*, 11:2543–2596, 2010.

L. Xiao and T. Zhang. A proximal stochastic gradient method with progressive variance reduction. *SIAM Journal on Optimization*, 24:2057–2075, 2014.

C.-H. Zhang and S. S. Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242, 2014.

T. Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the International Conference on Machine Learning*, 2004.

Appendix

A Proofs for results in Section 2

A.1 Proof of Lemma 2.3

Lemma A.1. *The following holds for any vectors:*

$$n^3(\|x_1\|_2^4 + \ldots + \|x_n\|_2^4) \geq \|x_1 + \ldots + x_n\|_2^4$$

Proof. Apply Cauchy-Schwartz inequality twice:

$$n(\|x_1\|_2^4 + \ldots + \|x_n\|_2^4) \geq (\|x_1\|_2^2 + \ldots + \|x_n\|_2^2)^2$$

and

$$n(\|x_1\|_2^2 + \ldots + \|x_n\|_2^2) \geq (\|x_1\|_2 + \ldots + \|x_n\|_2)^2.$$

Note that $\|x_1\|_2 + \ldots + \|x_n\|_2 \geq \|x_1 + \ldots + x_n\|_2$. □

Proof. The first item of Assumption 2.2 holds by definition. For notational simplicity, define

$$L(\Delta_{n-1}, \zeta_n) = (\nabla F(x_{n-1}) - \nabla F(x^*)) - (f(x_{n-1}, \zeta_n) - \nabla f(x^*, \zeta_n)),$$

By the definition of ξ_n and notice that $\nabla F(x^*) = 0$, we have,

$$\xi_n = L(\Delta_{n-1}, \zeta_n) - \nabla f(x^*, \zeta_n). \tag{45}$$

To prove Lemma 2.3, we need the fourth moment bounds on both $L(\Delta_{n-1}, \zeta_n)$ and $\nabla f(x^*, \zeta_n)$. By our assumption, $\nabla f(x^*, \zeta_n)$ has the bounded fourth moment. Now we upper bound $\mathbb{E}_{n-1}\|L(\Delta_{n-1}, \zeta_n)\|_2^m$ for $m = 2$ and $m = 4$. By the mean-value theorem and our assumption in (16), for $m = 2$ and 4, we have

$$\begin{aligned}\mathbb{E}_{n-1}\|\nabla f(x_{n-1}, \zeta_n) - \nabla f(x^*, \zeta_n)\|_2^m \leq & \mathbb{E}_{n-1}\left[\sup_x \left\|\nabla^2 f(x, \zeta_n)\right\|^m\right] \|\Delta_{n-1}\|_2^m \\ \leq & \|\Delta_{n-1}\|_2^m \mathbb{E}H(\zeta)^m.\end{aligned}$$

By the convexity of $\|x\|_2^m$ and Jensen's inequality,

$$\begin{aligned}\|\nabla F(x_{n-1}) - \nabla F(x^*)\|_2^m &= \|\mathbb{E}_{n-1}(\nabla f(x_{n-1}, \zeta_n) - \nabla f(x^*, \zeta_n))\|_2^m \\ &\leq \mathbb{E}_{n-1}\|\nabla f(x_{n-1}, \zeta_n) - \nabla f(x^*, \zeta_n)\|_2^m \\ &\leq \|\Delta_{n-1}\|_2^m \mathbb{E}H(\zeta)^m.\end{aligned}$$

By Lemma A.1, for $m = 2$ or 4,

$$\mathbb{E}_{n-1}\|L(\Delta_{n-1}, \zeta_n)\|_2^m \leq 2^m \|\Delta_{n-1}\|_2^m \mathbb{E}H(\zeta)^m.$$

Therefore, by the decomposition of ξ_n in (45),

$$\begin{aligned}\mathbb{E}_{n-1}\xi_n\xi_n^T = S - \mathbb{E}_{n-1}\nabla f(x^*, \zeta_n)L(\Delta_{n-1}, \zeta_n)^T - \mathbb{E}_{n-1}L(\Delta_{n-1}, \zeta_n)\nabla f(x^*, \zeta_n)^T \\ + \mathbb{E}_{n-1}L(\Delta_{n-1}, \zeta_n)L(\Delta_{n-1}, \zeta_n)^T,\end{aligned}$$

By Jensen's and Cauchy-Schwartz inequality,

$$\begin{aligned}&\left\|\mathbb{E}_{n-1}\nabla f(x^*, \zeta_n)L(\Delta_{n-1}, \zeta_n)^T\right\| \\ &\leq \mathbb{E}_{n-1}\left\|\nabla f(x^*, \zeta_n)L(\Delta_{n-1}, \zeta_n)^T\right\| \\ &= \mathbb{E}_{n-1}\|\nabla f(x^*, \zeta_n)\|_2 \|L(\Delta_{n-1}, \zeta_n)^T\|_2 \\ &\leq [\mathbb{E}_{n-1}\|\nabla f(x^*, \zeta_n)\|_2^2 \mathbb{E}_{n-1}\|L(\Delta_{n-1}, \zeta_n)\|_2^2]^{\frac{1}{2}} \\ &\leq 2\sqrt{\mathbb{E}_{n-1}\|\nabla f(x^*, \zeta_n)\|_2^2 \mathbb{E}H(\zeta)^2}\|\Delta_{n-1}\|_2\end{aligned}$$

Similarly,

$$\begin{aligned}&\left\|\mathbb{E}_{n-1}L(\Delta_{n-1}, \zeta_n)L(\Delta_{n-1}, \zeta_n)^T\right\| \\ &\leq \mathbb{E}_{n-1}\left\|L(\Delta_{n-1}, \zeta_n)L(\Delta_{n-1}, \zeta_n)^T\right\| = \mathbb{E}_{n-1}\|L(\Delta_{n-1}, \zeta_n)\|_2^2 \\ &\leq 4\mathbb{E}H(\zeta)^2\|\Delta_{n-1}\|_2^2\end{aligned}$$

Combining the above two inequalities, we arrive at item two of Assumption 2.2 with $\kappa = 1$.

Using the decomposition $\xi_n = L(\Delta_{n-1}, \zeta_n) - \nabla f(x^*, \zeta_n)$ and the fourth moment bounds on $L(\Delta_{n-1}, \zeta_n)$ and $\nabla f(x^*, \zeta_n)$, applying Lemma A.1 we get the third item. □

A.2 Lemma A.2

Lemma A.2. *Let a be a random variable on $\mathbb{R}^d$ with strictly positive density with respect to Lebesgue, and f be a continuous strictly positive function on $\mathbb{R}^d$. Then*

$$A := \mathbb{E}f(a)aa^T$$

is a PSD matrix with the minimum eigenvalue strictly above zero.

Proof. Let x be a unit-norm eigenvector of A associated with the minimum eigenvalue of A. It suffices to show that $x^T Ax > 0$. Let B be a neighborhood of x such that for any

$y \in B$, $\langle y, x \rangle > \frac{1}{2}$, and $f(y) > m$ for a constant $m > 0$. Then if we denote the density of a as $\mu(dy)$, we have

$$x^T A x = \int \mu(dy) f(y) x y y^T x \geq \int_B \mu(dy) f(y) x y y^T x \geq \frac{1}{4} \mu(B) m > 0.$$

□

A.3 Proof of Lemma 2.6

To prove Lemma 2.6, we need a simple lemma.

Lemma A.3. *Let z_n be a sequence in $\mathbb{R}^+$ that satisfies the recursion*

$$z_n \leq (1 - \lambda \eta_n) z_{n-1} + C \eta_n^{2+\beta}, \quad C, \lambda, \beta > 0$$

for a decreasing sequence η_n where $0 < \lambda \eta_n \leq 1$. Then for any $m \leq n - 1$,

$$z_n \leq \exp(-\lambda(n - m)\eta_n) z_m + C \eta_m^{1+\beta} \lambda^{-1}.$$

Proof. Using the recursion, we can find that

$$z_n \leq \prod_{k=m+1}^{n} (1 - \lambda \eta_k) z_m + C \sum_{k=m+1}^{n} \prod_{j=k+1}^{n} (1 - \lambda \eta_j) \eta_k^2.$$

Note that

$$\lambda \sum_{k=m+1}^{n} \prod_{j=k+1}^{n} (1 - \lambda \eta_j) \eta_k = 1 - \prod_{j=m+1}^{n} (1 - \lambda \eta_j) \leq 1.$$

Since $\eta_k \leq \eta_m$ for all $k \geq m + 1$,

$$C \sum_{k=m+1}^{n} \prod_{j=k+1}^{n} (1 - \lambda \eta_j) \eta_k^{2+\beta} \leq C \lambda^{-1} \eta_m^{1+\beta} \lambda \sum_{k=m+1}^{n} \prod_{j=k+1}^{n} (1 - \lambda \eta_j) \eta_k \leq C \lambda^{-1} \eta_m^{1+\beta}.$$

On the other hand, since $1 - x \leq \exp(-x)$ for $x \in [0, 1]$, so

$$\prod_{k=m+1}^{n} (1 - \lambda \eta_k) \leq \exp\left(-\lambda \sum_{k=m+1}^{n} \eta_k \right) \leq \exp(-\lambda(n - m)\eta_n).$$

Putting them back to the first inequality of this proof, we obtain our claim. □

With Lemma A.3 in place, we are ready to prove Lemma 2.6.

Proof of Lemma 2.6. From the recursion equation (9), we find that

$$\Delta_n = \Delta_{n-1} - \eta_n \nabla \widetilde{F}(\Delta_{n-1}) + \eta_n \xi_n,$$

where the shifted averaged loss function is $\widetilde{F}(\Delta) = F(\Delta + x^*)$. Therefore

$$\|\Delta_n\|_2^2 = \|\Delta_{n-1}\|_2^2 - 2\eta_n \langle \nabla \widetilde{F}(\Delta_{n-1}), \Delta_{n-1}\rangle + 2\eta_n \langle \xi_n, \Delta_{n-1}\rangle + \eta_n^2 \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^2. \quad (46)$$

The square of (46) is

$$\begin{aligned}
&\|\Delta_n\|_2^4 \\
=&\|\Delta_{n-1}\|_2^4 - 4\eta_n \langle \nabla \widetilde{F}(\Delta_{n-1}), \Delta_{n-1}\rangle \|\Delta_{n-1}\|_2^2 + 4\eta_n \langle \xi_n, \Delta_{n-1}\rangle \|\Delta_{n-1}\|_2^2 \\
&+4\eta_n^2 \langle \nabla \widetilde{F}(\Delta_{n-1}) - \xi_n, \Delta_{n-1}\rangle^2 + 2\eta_n^2 \|\Delta_{n-1}\|^2 \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^2 \\
&-4\eta_n^3 \langle \nabla \widetilde{F}(\Delta_{n-1}) - \xi_n, \Delta_{n-1}\rangle \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|^2 + \eta_n^4 \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^4.
\end{aligned}$$

By the property of the inner product, the fourth term above is less than twice of the fifth, which can be bounded by the following using Young's inequality,

$$2\eta_n^2 \|\Delta_{n-1}\|_2^2 \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^2 \le \frac{1}{3}\mu\eta_n \|\Delta_{n-1}\|_2^4 + 3\mu^{-1}\eta_n^3 \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^4;$$

Moreover, Hölder's inequality gives the following bound for the sixth term:

$$-4\eta_n^3 \langle \nabla \widetilde{F}(\Delta_{n-1}) - \xi_n, \Delta_{n-1}\rangle \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^2 \le \eta_n^3 \|\Delta_{n-1}\|_2^4 + 3\eta_n^3 \|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^4.$$

Combining all the inequality above back to the expansion of $\|\Delta_n\|_2^4$, we obtain

$$\begin{aligned}
\|\Delta_n\|_2^4 \le&\|\Delta_{n-1}\|_2^4 - 4\eta_n \langle \nabla \widetilde{F}(\Delta_{n-1}), \Delta_{n-1}\rangle \|\Delta_{n-1}\|_2^2 + 4\eta_n \langle \xi_n, \Delta_{n-1}\rangle \|\Delta_{n-1}\|_2^2 \\
&+ \mu\eta_n \|\Delta_{n-1}\|_2^4 + \eta_n^3 \|\Delta_{n-1}\|_2^4 + \eta_n^3 (9\mu^{-1} + \eta_n + 3)\|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^4. \quad (47)
\end{aligned}$$

Then using strong convexity of $\widetilde{F}$,

$$\langle \nabla \widetilde{F}(\Delta_{n-1}), \Delta_{n-1}\rangle \ge \widetilde{F}(\Delta_{n-1}) + \frac{\mu}{2}\|\Delta_{n-1}\|_2^2 \ge \frac{\mu}{2}\|\Delta_{n-1}\|_2^2.$$

By the fact that ξ_n is a martingale difference,

$$\mathbb{E}_{n-1} \langle \xi_n, \Delta_{n-1}\rangle \|\Delta_{n-1}\|_2^2 = 0.$$

Applying Young's inequality to Assumption 2.2 over ξ_n, using the Lipschitzness of ΔF in Assumption 2.1

$$\begin{aligned}
\mathbb{E}_{n-1}\|\xi_n - \nabla \widetilde{F}(\Delta_{n-1})\|_2^2 &\le 2\mathbb{E}_{n-1}\|\xi_n\|_2^2 + 2\|\nabla \widetilde{F}(\Delta_{n-1})\|_2^2 \\
&\le 2\mathrm{tr}(S) + 2\Sigma_2(\|\Delta_{n-1}\|_2^2 + \|\Delta_{n-1}\|_2^\kappa) + 2L_F^2 \|\Delta_{n-1}\|_2^2,
\end{aligned}$$

and likewise with Lemma A.1

$$\begin{aligned}\mathbb{E}_{n-1}\|\xi_n - \nabla\widetilde{F}(\Delta_{n-1})\|_2^4 &\leq 4\mathbb{E}_{n-1}\|\xi_n\|_2^4 + 4\|\nabla\widetilde{F}(\Delta_{n-1})\|_2^4 \\ &\leq 4\Sigma_4 + 4(\Sigma_4 + L_F^4)\|\Delta_{n-1}\|_2^4.\end{aligned}$$

Replacing propers terms in (46) and (47) with upper bounds above, we can find a constant C_1,

$$\mathbb{E}_{n-1}\|\Delta_n\|_2^2 \leq (1 - \mu\eta_n + C_1\eta_n^2)\|\Delta_{n-1}\|_2^2 + C_1\eta_n^2. \tag{48}$$

$$\mathbb{E}_{n-1}\|\Delta_n\|_2^4 \leq (1 - \mu\eta_n + C_1(\eta_n^3 + \eta_n^4))\|\Delta_{n-1}\|_2^4 + C_1\eta_n^3(\eta_n + 1). \tag{49}$$

Let

$$n_0 = \min\left\{n : C_1\eta_n^2 \leq \frac{1}{2}\mu\eta_n, C_1\eta_n^2(\eta_n + 1) \leq \frac{1}{2}\mu\eta_n, \mu\eta n(1-\alpha) \geq 8\alpha\log n\right\},$$

then for $n \geq n_0$,

$$\begin{aligned}\mathbb{E}_{n-1}\|\Delta_n\|_2^2 &\leq \left(1 - \frac{1}{2}\mu\eta_n\right)\|\Delta_{n-1}\|_2^2 + C_2\eta_n^2, \\ \mathbb{E}_{n-1}\|\Delta_n\|_2^4 &\leq \left(1 - \frac{1}{2}\mu\eta_n\right)\|\Delta_{n-1}\|_2^4 + C_2\eta_n^3.\end{aligned}$$

Then Lemma A.3 produces claim 1) for the second and forth moment, while we use Cauchy-Schwartz inequality it gives the bound for the first moment, as

$$\mathbb{E}_m\|\Delta_n\|_2 \leq \sqrt{\mathbb{E}_m\|\Delta_n\|_2^2}.$$

To see the claim 2), note that

$$\max\{C_1\eta_i^2, C_1(\eta_i^3 + \eta_i^4), C_1\eta_i^3(\eta_i + 1)\} \leq C_1\eta_i^2(\eta_i + 1)^2,$$

therefore applying discrete Gronwall's inequality to (48) and (49), we find that

$$\mathbb{E}\|\Delta_{n_0}\|_2^2 \leq C_3(1 + \|\Delta_0\|_2^2), \quad \mathbb{E}\|\Delta_{n_0}\|_2^4 \leq C_3(1 + \|\Delta_0\|_2^4),$$

where

$$C_3 = \left(\prod_{i=1}^{n_0}(1 + C_1\eta_i^2(\eta_i + 1)^2)\right)\left(C_1\sum_{i=1}^{n_0}\eta_i^2(\eta_i + 1)^2\right).$$

Then for $n \geq 2n_0$, using the claims from 1), there is a constant C_4 such that

$$\mathbb{E}_{n_0}\|\Delta_{\frac{n}{2}}\|_2^2 \leq \|\Delta_{n_0}\|_2^2 + C_4, \quad \mathbb{E}_{n_0}\|\Delta_{\frac{n}{2}}\|_2^4 \leq \|\Delta_{n_0}\|_2^4 + C_4,$$

and

$$\mathbb{E}_{\frac{n}{2}}\|\Delta_n\|_2^2 \leq \exp\left(-\tfrac{1}{4}n\mu\eta_n\right)\|\Delta_{\frac{n}{2}}\|_2^2 + C_4(\tfrac{n}{2})^{-\alpha},$$
$$\mathbb{E}_{\frac{n}{2}}\|\Delta_n\|_2^4 \leq \exp\left(-\tfrac{1}{2}n\mu\eta_n\right)\|\Delta_{\frac{n}{2}}\|_2^4 + C_4(\tfrac{n}{2})^{-2\alpha}.$$

So using the law of iterated expectation, and that

$$\exp(-\tfrac{1}{4}n\mu\eta_n) \leq \exp(-\tfrac{1}{4}\mu\eta n^{1-\alpha}) \leq n^{-2\alpha} \leq n^{-\alpha}$$

there is a constant C_5 that

$$\mathbb{E}\|\Delta_n\|_2^2 \leq n^{-\alpha}C_5\left(\|\Delta_0\|_2^2+1\right), \quad \mathbb{E}\|\Delta_n\|_2^4 \leq n^{-2\alpha}C_5\left(\|\Delta_0\|_2^4+1\right).$$

Applying Cauchy-Schwartz inequality again yields the claim for the first moment. □

Before we move on, let us apply the same arguments and prove a similar simpler result for the linear oracle sequence U_n.

Lemma A.4. *For the linear oracle sequence U_n defined by (31), it satisfies the following:*

$$\mathbb{E}\|U_n\|_2^2 \lesssim n^{-\alpha}(1+\|U_0\|_2^2+\|\Delta_0\|_2^2).$$

Proof. Clearly, we have the following recursion

$$\|U_n\|_2^2 = \|U_{n-1}\|_2^2 - 2\eta_n\langle AU_{n-1}, U_{n-1}\rangle + 2\eta_n\langle \xi_n, U_{n-1}\rangle + \eta_n^2\|\xi_n - AU_{n-1}\|_2^2.$$

Note that

$$\langle AU_{n-1}, U_{n-1}\rangle \geq \frac{\mu}{2}\|U_{n-1}\|_2^2,$$

and also that ξ_n is a martingale difference, so

$$\mathbb{E}_{n-1}\langle \xi_n, U_{n-1}\rangle = 0.$$

Moreover by Assumption 2.2, there is a constant C_1 such that

$$\begin{aligned}\mathbb{E}_{n-1}\|\xi_n - AU_{n-1}\|_2^2 &\leq 2\mathbb{E}_{n-1}\|\xi_n\|_2^2 + 2\|A\|^2\|U_{n-1}\|_2^2 \\ &\leq C_1(\|\Delta_{n-1}\|_2^2 + \|U_{n-1}\|_2^2 + 1).\end{aligned}$$

We find

$$\mathbb{E}_{n-1}\|U_n\|_2^2 \leq (1-\mu\eta_n + C_1\eta_n^2)\|U_{n-1}\|_2^2 + C_1(\eta_n^2 + \|\Delta_{n-1}\|_2^2).$$

Summing this inequality with (48) yields the following with $z_n := \|U_n\|_2^2 + \|\Delta_n\|_2^2$ and $C_2 = 2C_1$

$$\mathbb{E}_{n-1}z_n \leq (1-\mu\eta_n + C_2\eta_n^2)z_{n-1} + C_2\eta_n^2.$$

Then redoing the analysis in the proof of Lemma 2.6 after (48) yields that

$$\mathbb{E}\|U_n\|_2^2 \le \mathbb{E}z_n \lesssim n^{-\alpha}(1 + \|\Delta_0\|_2^2 + \|U_0\|_2^2).$$

□

B Consistency Proofs of the Plug-in Estimators

B.1 Proof of Lemma 3.2 for the consistency of A_n and S_n

Proof. We decompose $A_n - A$ into the following terms:

$$\begin{aligned} A_n - A &= \frac{1}{n}\sum_{i=1}^n \nabla^2 f(x_{i-1}, \zeta_i) - A \\ &= \left(\frac{1}{n}\sum_{i=1}^n \nabla^2 f(x^*, \zeta_i) - A\right) + \frac{1}{n}\sum_{i=1}^n \left(\nabla^2 f(x_{i-1}, \zeta_i) - \nabla^2 f(x^*, \zeta_i)\right). \end{aligned} \tag{50}$$

The first part can be bounded using the second moment bound, since by independence of ζ_i and ζ_j,

$$\mathbb{E}(\nabla^2 f(x^*, \zeta_i) - A)(\nabla^2 f(x^*, \zeta_j) - A)^T = \mathbb{1}_{i=j}[\mathbb{E}[\nabla^2 f(x^*, \zeta_i)][\nabla^2 f(x^*, \zeta_i)]^T - AA^T].$$

Therefore,

$$\begin{aligned} &\mathbb{E}\left(\frac{1}{n}\sum_{i=1}^n \nabla^2 f(x^*, \zeta_i) - A\right)\left(\frac{1}{n}\sum_{i=1}^n \nabla^2 f(x^*, \zeta_i) - A\right)^T \\ =&\frac{1}{n^2}\sum_{i=1}^n \mathbb{E}[\nabla^2 f(x^*, \zeta_i)][\nabla^2 f(x^*, \zeta_i)]^T - AA^T] \end{aligned}$$

which has its norm bounded by $\frac{1}{n}L_4$. Note that by Cauchy-Schwartz inequality, for any matrix B

$$\mathbb{E}\left\|BB^T\right\| = \mathbb{E}\left\|B\right\|^2 \ge [\mathbb{E}\left\|B\right\|]^2.$$

Therefore we have

$$\mathbb{E}\left\|\frac{1}{n}\sum_{i=1}^n \nabla^2 f(x^*, \zeta_i) - A\right\| \le \frac{\sqrt{L_4}}{\sqrt{n}}.$$

The second term in (50) can be bounded using Assumption 3.1,

$$\begin{aligned} &\left\|\mathbb{E}\frac{1}{n}\sum_{i=1}^n (\nabla^2 f(x_{i-1}, \zeta_i) - \nabla^2 f(x^*, \zeta_i))\right\| \\ \le&\frac{1}{n}\mathbb{E}\sum_{i=1}^n \left\|\nabla^2 f(x_{i-1}, \zeta_i) - \nabla^2 f(x^*, \zeta_i)\right\| \le \frac{L_2}{n}\sum_{i=1}^n \mathbb{E}\|x_i - x^*\|_2. \end{aligned}$$

Using Lemma 2.6 2), $\mathbb{E}\|x_i - x^*\| \lesssim i^{-\frac{\alpha}{2}}$. Note that by Lemma C.1 5) in below

$$\sum_{i=1}^{n} i^{-\frac{\alpha}{2}} \lesssim n^{1-\frac{\alpha}{2}}.$$

So the second term in (50) is of order $n^{-\frac{\alpha}{2}}$, which concludes the consistency proof for A_n.

For S_n, we decompose

$$\nabla f(x_{i-1}, \zeta_i) = \nabla f(x^*, \zeta_i) + (\nabla f(x_{i-1}, \zeta_i) - \nabla f(x^*, \zeta_i)) := X_i + Y_i,$$

then

$$S_n - S = \left(\frac{1}{n}\sum_{i=1}^{n} X_i X_i^T - S\right) + \frac{1}{n}\sum_{i=1}^{n} X_i Y_i^T + \frac{1}{n}\sum_{i=1}^{n} Y_i X_i^T + \frac{1}{n}\sum_{i=1}^{n} Y_i Y_i^T.$$

Note that

$$X_i = \nabla f(x^*, \zeta_i) = \nabla f(x^*, \zeta_i) - \nabla F(x^*)$$

which is ξ_n when $\Delta_{n-1} = 0$, so by part 3 of Assumption 2.2, we know the fourth moment of X_i is bounded. Then by $\mathbb{E}X_i X_i^T = S$ and the independence between the X_i's,

$$\mathbb{E}\left\|\frac{1}{n}\sum_{i=1}^{n} X_i X_i^T - S\right\|^2 \le \mathrm{tr}\left(\mathbb{E}\left(\frac{1}{n}\sum_{i=1}^{n} X_i X_i^T - S\right)^2\right) = \frac{1}{n}\mathrm{tr}((X_i X_i^T)^2 - S^2).$$

By Cauchy-Schwartz inequality,

$$\mathbb{E}\left\|\frac{1}{n}\sum_{i=1}^{n} X_i X_i^T - S\right\| \lesssim n^{-\frac{1}{2}}.$$

Further, by Assumption 3.1 and Lemma 2.6,

$$\begin{aligned}
\mathbb{E}\left\|Y_i Y_i^T\right\| &\le L_2 \mathbb{E}\left\|\Delta_{i-1}\right\|^2 \lesssim L_2 i^{-\alpha} \\
\mathbb{E}\left\|X_i X_i^T\right\| &\lesssim 1 \\
\mathbb{E}\left\|X_i Y_i^T\right\| &\le [\mathbb{E}\left\|X_i\right\|^2 \mathbb{E}\left\|Y_i\right\|^2]^{\frac{1}{2}} \lesssim i^{-\frac{\alpha}{2}}.
\end{aligned}$$

Summing up all the terms, we obtain $\mathbb{E}\left\|S_n - S\right\| \lesssim n^{-\frac{\alpha}{2}}$. □

B.2 Proof of Theorem 3.1 for consistency of the plug-in estimator

Before we provide the proof the consistency result of plug-in estimator in Theorem 3.1, we first provide the following lemma on matrix perturbation inequality for the inverse of a matrix.

Lemma B.1 (Matrix perturbation inequality for the inverse)**.** *Let $B = A + E$, where A and B are assumed to be invertible, and $\left\|A^{-1}E\right\| < \frac{1}{2}$. We have*

$$\left\|B^{-1} - A^{-1}\right\| \leq 2 \|E\| \left\|A^{-1}\right\|^2 \tag{51}$$

Proof. Using [Chang, 2006, Eq 3], we have

$$B^{-1} - A^{-1} = -A^{-1}E(I + A^{-1}E)^{-1}A^{-1},$$

which implies that

$$\begin{aligned}
\left\|B^{-1} - A^{-1}\right\| &= \left\|A^{-1}E(I + A^{-1}E)^{-1}A^{-1}\right\| \\
&\leq \left\|A^{-1}\right\|^2 \|E\| \left\|(I + A^{-1}E)^{-1}\right\| \\
&\leq \left\|A^{-1}\right\|^2 \|E\| \frac{1}{\lambda_{\min}(I + A^{-1}E)} \\
&\leq \left\|A^{-1}\right\|^2 \|E\| \frac{1}{1 - \|A^{-1}E\|}
\end{aligned}$$

Here $\lambda_{\min}$ is the minimum eigenvalue, and the last inequality uses Weyl's inequality $\lambda_{\min}(A + B) \geq \lambda_{\min}(A) - \|B\|$. The final result in (51) follows the assumption that $\left\|A^{-1}E\right\| < \frac{1}{2}$. □

Proof. Let us define

$$E_A = \widetilde{A}_n - A, \quad F_A = \widetilde{A}_n^{-1} - A^{-1}, \quad F_S = S_n - S,$$

Note that by Lemma 3.2, $\mathbb{E}\|F_S\| \lesssim n^{-\frac{\alpha}{2}}$ while by Markov inequality,

$$\begin{aligned}
\mathbb{E}\|E_A\| &\leq \mathbb{E}\|A_n - A\| + \|A\| \mathbb{P}(A_n \prec \tfrac{\lambda_A}{2} I) \\
&\lesssim n^{-\alpha/2} + \mathbb{P}(\|A_n - A\| \geq \tfrac{\lambda_A}{2}) \\
&\leq n^{-\alpha/2} + 2\lambda_A^{-1}\mathbb{E}\|A_n - A\| \lesssim n^{-\alpha/2}.
\end{aligned}$$

Then Markov inequality yields further that

$$\mathbb{P}(\left\|A^{-1}E_A\right\| \geq 1/2) \leq 2(\mathbb{E}\|E_A\|) \lesssim n^{-\alpha/2}.$$

Note that by construction, the following hold almost surely

$$\left\|\widetilde{A}_n^{-1}\right\| \leq 2\lambda_A^{-1}, \quad \left\|A^{-1}\right\| \leq \lambda_A^{-1}, \quad \|F_A\| \leq 3\lambda_A^{-1}.$$

By Lemma B.1,

$$\|F_A\| \leq \mathbb{1}_{\|A^{-1}E_A\| \leq 1/2} 2 \|E_A\| \left\|A^{-1}\right\|^2 + \mathbb{1}_{\|A^{-1}E_A\| \geq 1/2} 3\lambda_A^{-1},$$

therefore

$$\mathbb{E}\left\|F_A\right\| \leq 2\left\|A^{-1}\right\|^2 \mathbb{E}\left\|E_A\right\| + 3\lambda_A^{-1}\mathbb{P}(\left\|A^{-1}E_A\right\| \geq 1/2) \lesssim n^{-\frac{\alpha}{2}}.$$

Finally, we can decompose the error in the estimation into the following

$$\begin{aligned}\widetilde{A}_n^{-1}S_n\widetilde{A}_n^{-1} - A^{-1}SA^{-1} &= (A^{-1}+F_A)(S+F_S)(A^{-1}+F_A) - A^{-1}SA^{-1}, \\ &= (A^{-1}+F_A)F_S(A^{-1}+F_A) + A^{-1}SF_A + F_ASA^{-1} + F_ASF_A.\end{aligned}$$

Since each of the above terms involve either F_A or F_S, while all other terms are bounded by constants, the proof is completed. □

C Consistency proof of Batch-Means Estimator

C.1 Technical Lemmas

As a preparation for the consistency result, we introduce the following technical lemmas. Some variants of them have appeared in Polyak and Juditsky [1992] or Bach and Moulines [2011]. Since we need non-asymptotic bounds to investigate the relationship between the estimation error and M, N, we redo some key steps in Polyak and Juditsky [1992], Bach and Moulines [2011] with simpler notations.

Let us start with some approximations of various quantities related to the number of batches and batch sizes. These will simplify our discussion afterwards.

Lemma C.1. *Under the formulation of (28), the followings hold:*

1) The starting and ending indices of each non starting ($k \neq 0$) batches, and their neighbors $s_k - 1, s_k$ and $e_k, e_k + 1$ are of the same order,

$$s_k - 1 \asymp s_k \asymp e_k \asymp e_k + 1 \asymp (kN)^{\frac{1}{1-\alpha}}.$$

2) The total number of iterations used in our estimator can be estimated by:

$$S_M \asymp e_M \asymp (MN)^{\frac{1}{1-\alpha}}.$$

3) The number of iterations in each batch can be estimated by:

$$n_k = e_k - e_{k-1} \asymp k^{\frac{\alpha}{1-\alpha}} N^{\frac{1}{1-\alpha}}.$$

4) The sum of stepsize in each batch, which represents the log of decorrelation across that batch, is of order

$$\sum_{i=s_k}^{e_k} \eta_i \geq \eta_{e_k} n_k \gtrsim (kN)^{-\frac{\alpha}{1-\alpha}} k^{\frac{\alpha}{1-\alpha}} N^{\frac{1}{1-\alpha}} = N.$$

5) For any fixed $\gamma > 0$, the following estimate holds for sum of geometric sequence

$$M^{-1} \sum_{k=1}^{M} k^{-\gamma} \lesssim M^{-1} \int_1^{M+1} x^{-\gamma} dx \lesssim M^{(1-\gamma)^+ - 1}.$$

Proof. The first two inequalities comes from direct verifications, since $\frac{k+1}{k} \leq 2$ for $k \geq 1$. The third one comes from the fact that

$$e_k - e_{k-1} = [(k+1)^{\frac{1}{1-\alpha}} - k^{\frac{1}{1-\alpha}}] N^{\frac{1}{1-\alpha}} = \frac{N^{\frac{1}{1-\alpha}}}{1-\alpha} \int_k^{k+1} x^{\frac{\alpha}{1-\alpha}} dx,$$

and then observe that $k^{\frac{\alpha}{1-\alpha}} \leq x^{\frac{\alpha}{1-\alpha}} \leq 2^{\frac{\alpha}{1-\alpha}} k^{\frac{\alpha}{1-\alpha}}$. The fourth one is a corollary of the first and third. The last inequality comes from that $k+1 \leq 2k$ for $k \geq 1$ again. □

Recall the linear oracle sequence defined by (31): $U_n = (I - \eta_n A)\, \Delta_{n-1} + \eta_n \xi_n$. We now intruduce Lemma C.2 and provide the proof.

Lemma C.2. *For each positive integer j, let Y_j^k be a matrix sequence with $Y_j^j = I$ and for $k \geq j+1$*

$$Y_j^k = (I - \eta_k A) Y_j^{k-1} = \prod_{i=j+1}^{k} (I - \eta_i A),$$

where $\eta_k = \eta k^{-\alpha}$ with $\alpha \in [1/2, 1)$ and A is a PSD matrix with $\lambda_A > 0$ be any lower bound on the minimum eigenvalue of A. Then

1) For $i \in \{j, \ldots, k\}$, the following holds

$$\|Y_j^k\| \leq \exp\left(-\lambda_A \sum_{i=j+1}^{k} \eta_i\right) \leq \exp\left(-\lambda_A (k-j)\eta_k\right),$$

2) Let $S_j^k = \sum_{i=j+1}^{k} Y_j^i$ and $Z_j^k = \eta_j S_j^k - A^{-1}$, then

$$\left\|S_j^k\right\| \lesssim k^{\alpha}, \quad \left\|Z_j^k\right\| \lesssim k^{\alpha} j^{-1} + \exp\left(-\lambda_A \sum_{i=j}^{k} \eta_i\right).$$

3) When $s_1 \le j < k \le e_M$, for sufficiently large N, there exists a constant c such that

$$\left\|S_j^k\right\| \lesssim j^{\alpha}, \quad \left\|Z_j^k\right\| \lesssim j^{\alpha-1} + \exp\left(-c\lambda_A(k-j)\eta_j\right), \quad \left\|\eta_j S_j^k\right\| \lesssim 1,$$

as long as $M \le N^m$ for certain positive integer m.

Proof. Assume $Q\Lambda Q^T$ is the eigenvalue decomposition of A, then

$$\prod_{i=j+1}^{k}(I-\eta_i A) = Q\prod_{i=j+1}^{k}(I-\eta_i\Lambda)Q^T$$

which immediately gives us the claim 1) since $1-x \le \exp(-x)$ for $x \in [0,1]$. The first part of 2) comes as an immediate consequence since:

$$\|S_j^k\| \le \sum_{i=j+1}^{k}\|Y_j^i\| \le \sum_{i=j+1}^{k}\exp(-\lambda_A(i-j)\eta_k) \le \lambda_A^{-1}\eta_k^{-1},$$

while $\eta_k \asymp k^{-\alpha}$. To show the second part of 2), we sum the identity $Y_j^i = Y_j^{i-1} - \eta_k A Y_j^{i-1}$ from $i = j+1$ to $i = k+1$, which gives us:

$$Y_j^{k+1} = I - \sum_{i=j}^{k}\eta_k A Y_j^i.$$

This can be rewritten as

$$\sum_{i=j}^{k}\eta_i Y_j^i = A^{-1} - A^{-1}Y_j^{k+1}.$$

Therefore,

$$\begin{aligned}\left\|Z_j^k\right\| &\le \left\|\eta_i S_i^n - \sum_{i=j}^{k}\eta_i Y_j^i\right\| + \left\|\sum_{i=j}^{k}\eta_i Y_j^i - A^{-1}\right\| \\ &\le \left\|\eta_i S_i^n - \sum_{i=j}^{k}\eta_i Y_j^i\right\| + \left\|A^{-1}\right\|\exp\left(-\lambda_A\sum_{m=j+1}^{k+1}\eta_m\right).\end{aligned}$$

In order for us to have claim 2), it suffices to bound the first part. Notice that for $j \le m \le i$,

$$\begin{aligned}\eta_{m-1} - \eta_m = \eta((m-1)^{-\alpha} - m^{-\alpha}) &= \frac{\eta}{\alpha}\int_{m-1}^{m} x^{-\alpha-1}dx \\ &\le \frac{2^{\alpha+1}\eta}{\alpha}m^{-\alpha-1} \le 4(\eta\alpha)^{-1}\eta_m\eta_{i+1}(i+1)^{\alpha}j^{-1},\end{aligned}$$

therefore

$$\eta_j - \eta_i = \sum_{m=j+1}^{i} (\eta_{m-1} - \eta_m) \leq \sum_{m=j+1}^{i} 4(\eta\alpha)^{-1}\eta_m\eta_i(i+1)^\alpha j^{-1}$$
$$\leq 4(\eta\alpha)^{-1}(i+1)^\alpha j^{-1}\eta_{i+1}\sum_{m=j+1}^{i}\eta_m.$$

Using $S_j^k = \sum_{i=j+1}^{k} Y_j^i$,

$$\begin{aligned}
&\left\|\eta_j S_j^k - \sum_{i=j}^{k}\eta_i Y_j^i\right\| \\
&\leq \sum_{i=j}^{k}(\eta_j - \eta_i)\left\|Y_j^i\right\| + \eta_j \\
&\leq 4(\eta\alpha)^{-1}k^\alpha j^{-1}\sum_{i=j}^{k}\eta_{i+1}\Big(\sum_{m=j+1}^{i}\eta_m\Big)\exp\Big(-\lambda_A\sum_{m=j+1}^{i}\eta_m\Big) + \eta j^{-\alpha}
\end{aligned}$$

Note that $xe^{-\lambda_A x} \leq e^{\lambda_A} y e^{-\lambda_A y}$ when $x \leq y \leq x+1$, so

$$\begin{aligned}
&\eta_{i+1}\Big(\sum_{m=j+1}^{i}\eta_m\Big)\exp\Big(-\lambda_A\sum_{m=j+1}^{i}\eta_m\Big) \\
&= \int_{\sum_{m=j+1}^{i}\eta_m}^{\sum_{m=j+1}^{i+1}\eta_m}\Big(\sum_{m=j+1}^{i}\eta_m\Big)\exp\Big(-\lambda_A\sum_{m=j+1}^{i}\eta_m\Big) \\
&\leq e^{\lambda_A}\int_{\sum_{m=j+1}^{i}\eta_m}^{\sum_{m=j+1}^{i+1}\eta_m} y e^{-\lambda_A y}dx.
\end{aligned}$$

So

$$\sum_{i=j}^{k}\eta_{i+1}\Big(\sum_{m=j+1}^{i}\eta_m\Big)\exp\Big(-\lambda_A\sum_{m=j+1}^{i}\eta_m\Big) \leq e^{\lambda_A}\int_0^\infty x e^{-\lambda_A x}dx,$$

which is a constant. Note that $j^{-\alpha} \leq j^{\alpha-1} \leq k^\alpha j^{-1}$, therefore

$$\left\|Z_j^k\right\| \leq \left\|\eta_i S_i^n - \sum_{i=j}^{k}\eta_i Y_j^i\right\| + \left\|\sum_{i=j}^{k}\eta_i Y_j^i - A^{-1}\right\| \lesssim k^\alpha j^{-1} + \exp\Big(-\lambda_A\sum_{i=j+1}^{k}\eta_i\Big).$$

For claim 3), when j and k are in neighboring batches, the second claim already produces the results, since the k/j is bounded by $3^{\frac{1}{1-\alpha}}$. Otherwise there is a batch between j, k, that is for an m the following holds,

$$j \leq e_{m-1} \leq e_m \leq k.$$

Note that

$$S_j^k = S_j^{e_m} + Y_j^{e_m} S_{e_m}^k$$

so

$$Z_j^k = \eta_j S_j^k - A^{-1} = Z_j^{e_m} + Y_{e_m}^k S_{e_m}^k$$

Then because e_m and j are in neighboring batches, by Lemma C.1 claim 3, there is a c such that

$$\left\| Z_j^{e_m} \right\| \lesssim K_1 j^{\alpha-1} + \exp\left(-\lambda_A \sum_{i=j+1}^{e_m} \eta_i \right) \leq j^{\alpha-1} + \exp(-c\lambda_A N).$$

Because e_m is at least one batch away from j, so by the first claim of this lemma and Lemma C.1 4)

$$\left\| Y_j^{e_m} \right\| \leq \exp\left(-\lambda_A \sum_{i=s_m}^{e_m} \eta_i \right) \leq \exp(-\lambda_A c N),$$

while by the second claim,

$$\left\| S_{e_m}^k \right\| \lesssim e_M^\alpha \asymp (NM)^{\frac{\alpha}{1-\alpha}}.$$

Then because M is of polynomial order of N, so for N large enough,

$$\| Y_{e_m}^k S_{e_m}^k \| \lesssim \eta^\alpha (NM)^{\frac{\alpha}{1-\alpha}} \exp(-\lambda_A c N) \leq e_M^{\alpha-1} \leq j^{\alpha-1}$$

Therefore,

$$\left\| Z_j^k \right\| \lesssim j^{\alpha-1} + \exp(-c\lambda_A N) \asymp j^{\alpha-1}.$$

□

C.2 A linear framework

This section is devoted to the linear framework illustrated by Lemma 3.3. As we will see, the following two quantities

$$\widehat{S} := M^{-1} \sum_{k=1}^{M} n_k^{-1} \left(\sum_{i=s_k}^{e_k} \xi_i \right) \left(\sum_{i=s_k}^{e_k} \xi_i \right)^T$$

and $M^{-1} \sum_{k=1}^{M} n_k \overline{U}_{n_k} \overline{U}_{n_k}^T$ converge to the asymptotic covariance $A^{-1} S A^{-1}$. The strategy of the proof below is first showing $\widehat{S}$ converges to the right covariance, then $M^{-1} \sum_{k=1}^{M} n_k \overline{U}_{n_k} \overline{U}_{n_k}^T$ is not far from the $\widehat{S}$, and eventually reaches Lemma 3.3.

Lemma C.3. *Under Assumptions 2.1 and 2.2, the following*

$$\widehat{S} := M^{-1}\sum_{k=1}^{M} n_k^{-1}\left(\sum_{i=s_k}^{e_k}\xi_i\right)\left(\sum_{i=s_k}^{e_k}\xi_i\right)^T \tag{52}$$

is a consistent estimator of S, namely

$$\mathbb{E}\|\widehat{S}-S\|^2 \lesssim d\left(N^{-\frac{\kappa\alpha}{1-\alpha}}M^{-\beta_\alpha}+M^{-1}\right)$$

with $\beta_\alpha := 1-\left(1-\frac{\kappa\alpha}{2-2\alpha}\right)^+$.

Proof. Since

$$\mathbb{E}\left\|(\widehat{S}-S)^2\right\| \le \mathbb{E}\mathrm{tr}(\widehat{S}-S)^2 = \mathrm{tr}\mathbb{E}(\widehat{S}-S)^2 \le d\left\|\mathbb{E}(\widehat{S}-S)^2\right\|,$$

where $\mathrm{tr}(\cdot)$ denotes the trace of a matrix, it suffices to show $\|\mathbb{E}(\widehat{S}-S)^2\|$ follows the claimed bound. Notice that

$$\mathbb{E}\widehat{S} = M^{-1}\sum_{k=1}^{M} n_k^{-1}\sum_{i=s_k}^{e_k}\mathbb{E}\xi_i\xi_i^T = M^{-1}\sum_{k=1}^{M} n_k^{-1}\sum_{i=s_k}^{e_k}(S+\mathbb{E}\Sigma(\Delta_{i-1}))$$

So

$$\mathbb{E}(\widehat{S}-S)^2 = \mathbb{E}\widehat{S}^2 - S^2 + M^{-1}S\sum_{k=1}^{M} n_k^{-1}\sum_{i=s_k}^{e_k}\mathbb{E}\Sigma(\Delta_{i-1}) + M^{-1}\sum_{k=1}^{M} n_k^{-1}\sum_{i=s_k}^{e_k}\mathbb{E}\Sigma(\Delta_{i-1})S. \tag{53}$$

The third term can be bounded by the following

$$\begin{aligned}
&\left\|M^{-1}S\sum_{k=1}^{M} n_k^{-1}\sum_{i=s_k}^{e_k}\mathbb{E}\Sigma(\Delta_{i-1})\right\| \\
\le&\Sigma_2 M^{-1}\|S\|\sum_{k=1}^{M} n_k^{-1}\sum_{i=s_k}^{e_k}\mathbb{E}(\|\Delta_{i-1}\|_2^\kappa + \|\Delta_{i-1}\|_2^2) \\
\overset{(a)}{\lesssim}& M^{-1}\sum_{k=1}^{M} s_k^{-\kappa\alpha/2}(\|\Delta_0\|_2^\kappa + \|\Delta_0\|_2^2+1) \\
\overset{(b)}{\lesssim}& s_1^{-\kappa\alpha/2} \asymp N^{-\frac{\kappa\alpha}{2(1-\alpha)}},
\end{aligned}$$

where (a) uses Lemma 2.6 and (b) uses $s_k > s_1$. Likewise, we can bound the fourth term in (53) as well. For the computation of $\mathbb{E}\widehat{S}^2$, which is

$$\mathbb{E}\widehat{S}^2 = M^{-2}\sum_{j=1}^{M}\sum_{k=1}^{M} n_j^{-1}n_k^{-1}\sum_{i_1,i_2=s_j}^{e_j}\sum_{i_3,i_4=s_k}^{e_k}\mathbb{E}\xi_{i_1}\xi_{i_2}^T\xi_{i_3}\xi_{i_4}^T,$$

using the marginal difference nature of ξ_i, one can see that the terms are 0 unless indices come in pairs, i.e. $i_1 = i_2, i_3 = i_4$ or $i_1 = i_3, i_2 = i_4$ or $i_1 = i_4, i_2 = i_3$. Also notice that the latter two only happens when all the indices are in the same batch. So

$$\mathbb{E}\widehat{S}^2 = \frac{1}{M^2}\Bigg[\sum_{j=1}^{M}\sum_{k=1}^{M} n_j^{-1}n_k^{-1}\sum_{i_1=s_j}^{e_j}\sum_{i_3=s_k}^{e_k}\mathbb{E}\xi_{i_1}\xi_{i_1}^T\xi_{i_3}\xi_{i_3}^T + \sum_{j=1}^{M} n_j^{-2}\sum_{i_1,i_3=s_j,i_1\neq i_3}^{e_j}\mathbb{E}\xi_{i_1}\xi_{i_3}^T\xi_{i_1}\xi_{i_3}^T + \mathbb{E}\xi_{i_1}\xi_{i_3}^T\xi_{i_3}\xi_{i_1}^T\Bigg],$$

Then using that

$$S^2 = M^{-2}\sum_{j,k=1}^{M} n_j^{-1}n_k^{-1}\sum_{i_1=s_j}^{e_j}\sum_{i_3=s_k}^{e_k} S^2$$

we have

$$\begin{aligned}\left\|\mathbb{E}\widehat{S}^2 - S^2\right\| \leq & M^{-2}\sum_{j=1}^{M}\sum_{k=1}^{M} n_j^{-1}n_k^{-1}\sum_{i_1=s_j}^{e_j}\sum_{i_3=s_k}^{e_k}\left\|\mathbb{E}\xi_{i_1}\xi_{i_1}^T\xi_{i_3}\xi_{i_3}^T - S^2\right\| \\ & + M^{-2}\sum_{j=1}^{M} n_j^{-2}\sum_{i_1,i_2=s_j}^{e_j}\left\|\mathbb{E}\xi_{i_1}\xi_{i_2}^T\xi_{i_1}\xi_{i_2}^T\right\| + \left\|\mathbb{E}\xi_{i_1}\xi_{i_2}^T\xi_{i_2}\xi_{i_1}^T\right\|. \end{aligned} \tag{54}$$

First, we setup a upper bound of $\left\|\mathbb{E}\xi_{i_1}\xi_{i_1}^T\xi_{i_3}\xi_{i_3}^T - S^2\right\|$ when $|j-k|>1$, that is when i_1 and i_3 are from two non-neighboring batches. Without lost of generality, we assume that $i_3 > i_1$:

$$\begin{aligned}\left\|\mathbb{E}(\xi_{i_1}\xi_{i_1}^T)(\xi_{i_3}\xi_{i_3}^T) - S^2\right\| &= \left\|\mathbb{E}(\xi_{i_1}\xi_{i_1}^T)(S + \mathbb{E}_{i_1}\Sigma(\Delta_{i_3-1})) - S^2\right\| \\ &\overset{(a)}{=} \left\|S^2 + \mathbb{E}\Sigma(\Delta_{i_1-1})S + \mathbb{E}\xi_{i_1}\xi_{i_1}^T\mathbb{E}_{i_1}\Sigma(\Delta_{i_3-1}) - S^2\right\| \\ &\leq \Sigma_2\|S\|\,\mathbb{E}\|\Delta_{i_1-1}\|_2^{\kappa} + \left\|\mathbb{E}\xi_{i_1}\xi_{i_1}^T\mathbb{E}_{i_1}\Sigma(\Delta_{i_3-1})\right\|,\end{aligned}$$

where (a) uses Assumption 2.2. Also notice that,

$$\begin{aligned}\left\|\mathbb{E}\xi_{i_1}\xi_{i_1}^T\mathbb{E}_{i_1}\Sigma(\Delta_{i_3-1})\right\| &\leq \mathbb{E}\left\|\xi_{i_1}\xi_{i_1}^T\right\|\mathbb{E}_{i_1}\left\|\Sigma(\Delta_{i_3-1})\right\| \\ &\leq \Sigma_2\mathbb{E}\|\xi_{i_1}\|_2^2\mathbb{E}_{i_1}(\|\Delta_{i_3-1}\|_2^{\kappa} + \|\Delta_{i_3-1}\|_2^2) \\ &\overset{(a)}{\lesssim} \mathbb{E}\|\xi_{i_1}\|_2^2\left[\exp\left(-\lambda_A\sum_{k=i_1+1}^{i_3-1}\eta_k\right)\|\Delta_{i_1}\|_2^{\kappa} + i_1^{-\kappa\alpha/2}\right] \\ &\overset{(b)}{\leq} \mathbb{E}\|\xi_{i_1}\|_2^2\left[\exp(-\lambda_A cN)\|\Delta_{i_1}\|_2^{\kappa} + i_1^{-\kappa\alpha/2}\right] \\ &\overset{(c)}{\leq} i_1^{-\kappa\alpha/2}\left[\mathbb{E}\|\xi_{i_1}\|_2^2 + \sqrt{\mathbb{E}\|\xi_{i_1}\|_2^4\mathbb{E}\|\Delta_{i_1}\|_2^{2\kappa}}\right] \\ &\overset{(d)}{\lesssim} i_1^{-\kappa\alpha/2},\end{aligned}$$

where (a) uses the second claim in Lemma 2.6, (b) uses the fourth claim in Lemma C.1, (c) is due to the fact that $\exp(-c\lambda_A N) \ll s_M^{-\alpha}$ and (d) uses the second and third claim in Lemma 2.6.

Now using the second claim in Lemma 2.6, we obtain the following bound for i_1, i_3 from non-neighboring batches:

$$\left\|\mathbb{E}(\xi_{i_1}\xi_{i_1}^T)(\xi_{i_3}\xi_{i_3}^T) - S^2\right\| \lesssim i_1^{-\kappa\alpha/2} \leq i_1^{-\kappa\alpha/2} + i_3^{-\kappa\alpha/2}.$$

We added $i_3^{-\kappa\alpha/2}$ so the bound holds also for the case that $i_1 > i_3$. On the other hand, the following bound holds for general choice of $i_1, \ldots, i_4$ with a proper C_1 by Lemma 2.6 claim 3:

$$\mathbb{E}\left\|\xi_{i_1}\xi_{i_2}^T\xi_{i_3}\xi_{i_4}^T\right\| \leq \frac{1}{4}\mathbb{E}[\|\xi_{i_1}\|_2^4 + \|\xi_{i_2}\|_2^4 + \|\xi_{i_3}\|_2^4 + \|\xi_{i_4}\|_2^4] \leq C_1,$$

and also

$$\|\mathbb{E}\xi_{i_1}\xi_{i_2}^T\xi_{i_3}\xi_{i_4}^T - S^2\| \leq \mathbb{E}\|\xi_{i_1}\xi_{i_2}^T\xi_{i_3}\xi_{i_4}^T\| + \|S^2\| \leq C_1 + \|S^2\|.$$

Therefore the inequality (54) gives the following:

$$\begin{aligned}
\left\|\mathbb{E}\widehat{S}^2 - S^2\right\| \lesssim & M^{-2} \sum_{|j-k|>1} n_j^{-1} n_k^{-1} \sum_{i_1=s_j}^{e_j} \sum_{i_3=s_k}^{e_k} (i_1^{-\kappa\alpha/2} + i_3^{-\kappa\alpha/2}) \\
& + M^{-2} \sum_{|j-k|\leq 1} n_j^{-1} n_k^{-1} \sum_{i_1=s_j}^{e_j} \sum_{i_3=s_k}^{e_k} (C_1 + \left\|S^2\right\|) \\
& + 2M^{-2} \sum_{j=1}^{M} n_j^{-2} \sum_{i_1, i_1=s_j}^{e_j} C_1 \\
\lesssim & M^{-1} \left[\sum_{k=1}^{M} s_k^{-\kappa\alpha/2}\right] + 3(C_1 + \left\|S^2\right\|)M^{-1} \\
\lesssim & N^{-\frac{\kappa\alpha}{2-2\alpha}} M^{[1-\frac{\kappa\alpha}{2(1-\alpha)}]^+ - 1} + M^{-1},
\end{aligned}$$

where we used estimations from Lemma C.1 1) and 5). □

Lemma C.4. *Under Assumptions 2.1 and 2.2, the following holds for the linear oracle sequence*

$$\mathbb{E}\left\|M^{-1}\sum_{k=1}^{M} n_k \overline{U}_{n_k} \overline{U}_{n_k}^T - A^{-1}SA^{-1}\right\| \lesssim N^{-\frac{\kappa\alpha}{4-4\alpha}} M^{-\frac{\beta_\alpha}{2}} + M^{-\frac{1}{2}} + N^{-\frac{1}{2}},$$

with β_α given by Lemma C.3.

Proof. We define the following matrices sequences through recursion as in Lemma C.2,

$$Y_i^j = (I - \eta_j A) Y_i^{j-1}, \quad Y_i^i = I_d.$$

Then the recursion of U_n can be rewritten as

$$\begin{aligned} U_n &= Y_{n-1}^n U_{n-1} - Y_n^n \eta_n \xi_n \\ &= Y_{n-2}^n U_{n-2} - Y_{n-1}^n \eta_{n-1} \xi_{n-1} - Y_n^n \eta_n \xi_n \\ &= Y_{s-1}^n U_{s-1} - \sum_{i=s}^{n} Y_i^n \eta_i \xi_i. \end{aligned}$$

Therefore $\overline{U}_{n_k}$ can be written as

$$\begin{aligned} \overline{U}_{n_k} &= n_k^{-1} \left[\sum_{l=s_k}^{e_k} Y_{s_k-1}^l U_{s_k-1} - \sum_{l=s_k}^{e_k} \sum_{i=s_k}^{l} Y_i^l \eta_i \xi_i \right] \\ &= n_k^{-1} \left[\sum_{l=s_k}^{e_k} Y_{s_k-1}^l U_{s_k-1} - \sum_{i=s_k}^{e_k} \left(\sum_{l=i}^{e_k} Y_i^l \right) \eta_i \xi_i \right] \\ &= n_k^{-1} \left[\sum_{l=s_k}^{e_k} Y_{s_k-1}^l U_{s_k-1} - \sum_{i=s_k}^{e_k} \left(S_i^{e_k} + I \right) \eta_i \xi_i \right] \\ &= n_k^{-1} S_{s_k-1}^{e_k} U_{s_k-1} - n_k^{-1} A^{-1} \sum_{i=s_k}^{e_k} \xi_i - n_k^{-1} \sum_{i=s_k}^{e_k} Z_i^{e_k} \xi_i - n_k^{-1} \sum_{i=s_k}^{e_k} \eta_i \xi_i \end{aligned}$$

Here we denoted

$$S_m^{e_k} = \sum_{l=m+1}^{e_k} Y_m^l = \sum_{l=m}^{e_k} Y_m^l - I, \quad Z_m^{e_k} = \eta_m S_m^{e_k} - A^{-1},$$

as in Lemma C.2. Denote

$$A_k := -A^{-1} \sum_{i=s_k}^{e_k} \xi_i, \quad B_k := S_{s_k-1}^{e_k} U_{s_k-1} - \sum_{i=s_k}^{e_k} Z_i^{e_k} \xi_i - \sum_{i=s_k}^{e_k} \eta_i \xi_i,$$

then

$$\begin{aligned} M^{-1} \sum_{k=1}^{M} n_k \overline{U}_{n_k} \overline{U}_{n_k}^T =& M^{-1} \sum_{k=1}^{M} n_k^{-1} A_k A_k^T \\ &+ M^{-1} \sum_{k=1}^{M} n_k^{-1} \left[A_k B_k^T + B_k A_k^T + B_k B_k^T \right]. \end{aligned} \tag{55}$$

Denote

$$\widehat{S} := M^{-1}\sum_{k=1}^{M} n_k^{-1}\left(\sum_{i=s_k}^{e_k}\xi_i\right)\left(\sum_{i=s_k}^{e_k}\xi_i\right)^T.$$

By Lemma C.3, we have for sufficiently large N

$$\begin{aligned}\mathbb{E}\left\|M^{-1}\sum_{k=1}^{M} n_k^{-1}A_kA_k^T - A^{-1}SA^{-1}\right\|^2 =&\mathbb{E}\left\|A^{-1}(\widehat{S}-S)A^{-1}\right\|^2\\ &\lesssim \left[N^{-\frac{\kappa\alpha}{2(1-\alpha)}}M^{-\beta_\alpha}+M^{-1}\right].\end{aligned}$$

By Cauchy-Schwartz inequlaity,

$$\mathbb{E}\left\|M^{-1}\sum_{k=1}^{M} n_k^{-1}A_kA_k^T - A^{-1}SA^{-1}\right\| \lesssim N^{-\frac{\kappa\alpha}{4-4\alpha}}M^{-\frac{\beta_\alpha}{2}}+M^{-\frac{1}{2}}.$$

Next, note that

$$\mathbb{E}\|B_k\|_2^2 = \mathbb{E}\left\|B_kB_k^T\right\| \le \mathbb{E}\mathrm{tr}[B_kB_k^T] \le d\left\|\mathbb{E}B_kB_k^T\right\|.$$

Using that ξ_n are martingale differences

$$\begin{aligned}&\mathbb{E}B_kB_k^T\\ &= \mathbb{E}\left[S^{e_k}_{s_k-1}U_{s_k-1}-\sum_{i=s_k}^{e_k}Z_i^{e_k}\xi_i-\sum_{i=s_k}^{e_k}\eta_i\xi_i\right]\left[S^{e_k}_{s_k-1}U_{s_k-1}-\sum_{i=s_k}^{e_k}Z_i^{e_k}\xi_i-\sum_{i=s_k}^{e_k}\eta_i\xi_i\right]^T,\\ &= S^{e_k}_{s_k-1}\mathbb{E}\left[U_{s_k-1}U^T_{s_k-1}\right](S^{e_k}_{s_k-1})^T+\sum_{i=s_k}^{e_k}(Z_i^{e_k}+\eta_iI)\mathbb{E}(\xi_i\xi_i^T)(Z_i^{e_k}+\eta_iI)^T,\end{aligned}$$

From C.2 claim 3,

$$\mathrm{tr}(\mathbb{E}B_kB_k^T) \le \left\|S^{e_k}_{s_k-1}\right\|^2\mathbb{E}\|U_{s_k-1}\|_2^2+\sum_{i=s_k}^{e_k}\|Z_i^{e_k}+\eta_iI\|^2\mathbb{E}\|\xi_i\|_2^2,$$

Note that from Lemma C.2 claim 3

$$\|Z_i^{e_k}+\eta_iI\| \le \|Z_i^{e_k}\|+\eta i^{-\alpha} \lesssim i^{\alpha-1}+\exp\left(-c\lambda_A(e^k-i)\eta_i\right)+i^{-\alpha}.$$

Using Young's inequality and the estimate from Lemma C.1, and $i^{-\alpha} \le i^{\alpha-1}$ with $\alpha \in [1/2,1)$,

$$\|Z_i^{e_k}+\eta_iI\|^2 \lesssim {s_k}^{2\alpha-2}+\exp(-2c\lambda_A(e_k-i)\eta_{e_k}).$$

Thus from the bounds of $\|U_n\|_2^2$,

$$\mathrm{tr}(\mathbb{E}B_kB_k^T) \lesssim s_k^\alpha+\sum_{i=s_k}^{e_k}\|Z_i^{e_k}+\eta_iI\|^2 \lesssim s_k^\alpha+n_k{s_k}^{2\alpha-2}+\sum_{i=s_k}^{e_k}\exp(-2c\lambda_A(e_k-i)\eta_{e_k})$$

Note that

$$\sum_{i=s_k}^{e_k} \exp(-2c\lambda_A(e_k-i)\eta_{e_k}) \le \sum_{i=0}^{\infty} \exp(-2c\lambda_A i\eta_{e_k}) = [2c\lambda_A\eta_{e_k}]^{-1}.$$

Then based on Lemma C.1,

$$n_k s_k^{2\alpha-2} + \eta_{e_k}^{-1} \lesssim N^{-\frac{1-2\alpha}{1-\alpha}} k^{-\frac{2-\alpha}{1-\alpha}} + N^{-\frac{\alpha}{1-\alpha}} k^{-\frac{\alpha}{1-\alpha}} \lesssim s_k^{\alpha}.$$

Using the bound that $e_k^{\alpha}(s_k-1)^{-\alpha} \le 4$,

$$M^{-1}\sum_{k=1}^{M} n_k^{-1}\mathbb{E}\|B_k\|_2^2 \lesssim M^{-1}\sum_{k=1}^{M} n_k^{-1}s_k^{\alpha}.$$

By Lemma C.1,

$$n_k^{-1}s_k^{\alpha} \lesssim N^{-\frac{1}{1-\alpha}} k^{-\frac{\alpha}{1-\alpha}} N^{\frac{\alpha}{1-\alpha}} k^{\frac{\alpha}{1-\alpha}} = N^{-1}.$$

So

$$M^{-1}\sum_{k=1}^{M} n_k^{-1}\mathbb{E}\|B_k\|_2^2 \lesssim dN^{-1}.$$

On the other hand, by Lemma 2.6,

$$\begin{aligned} n_k^{-1}\mathbb{E}\|A_k\|_2^2 \le & n_k^{-1} d\,\mathrm{tr}\mathbb{E}A_k A_k^T \\ \le & n_k^{-1} d\|A^{-1}\|^2 \left[n_k \mathrm{tr}(S) + 2\Sigma_2 \sum_{i=s_k}^{e_k} \mathbb{E}(\|\Delta_i\|_2^{2\kappa} + \|\Delta_i\|_2^4) \right] \\ \lesssim & d(1+\|\Delta_0\|_2^2). \end{aligned}$$

So

$$n_k^{-1}\mathbb{E}\|A_k\|_2\|B_k\|_2 \lesssim \sqrt{n_k^{-1}\mathbb{E}\|A_k\|_2^2}\sqrt{n_k^{-1}\mathbb{E}\|B_k\|_2^2} \lesssim dN^{-\frac{1}{2}}.$$

Therefore, in conclusion,

$$\begin{aligned} & \mathbb{E}\left\| M^{-1}\sum_{k=1}^{M} \overline{U}_{n_k}\overline{U}_{n_k}^T - A^{-1}SA^{-1} \right\| \\ \le & \mathbb{E}\left\| M^{-1}\sum_{k=1}^{M} n_k^{-1}A_k A_k^T - A^{-1}SA^{-1} \right\| \\ & + 2M^{-1}\sum_{k=1}^{M} n_k^{-1}(\mathbb{E}\|A_k\|_2\|B_k\|_2 + \mathbb{E}\|B_k\|_2^2) \\ \lesssim & N^{-\frac{\kappa\alpha}{4-4\alpha}} M^{-\frac{\beta\alpha}{2}} + M^{-\frac{1}{2}} + N^{-\frac{1}{2}}. \end{aligned}$$

□

C.3 Proof of Lemma 3.3

Proof. Using the identity $\sum_{k=1}^{M} n_k \overline{U}_M = \sum_{k=1}^{M} n_k \overline{U}_{n_k}$, we can easily find

$$\begin{aligned} &M^{-1}\sum_{k=1}^{M} n_k(\overline{U}_{n_k} - \overline{U}_M)(\overline{U}_{n_k} - \overline{U}_M)^T \\ =&M^{-1}\sum_{k=1}^{M} n_k \overline{U}_{n_k}\overline{U}_{n_k}^T - M^{-1}\sum_{k=1}^{M} n_k \overline{U}_M \overline{U}_M^T. \end{aligned}$$

Lemma C.4 already shows that

$$\mathbb{E}\left\| M^{-1}\sum_{k=1}^{M} n_k \overline{U}_{n_k}\overline{U}_{n_k}^T - A^{-1}SA^{-1}\right\| \lesssim N^{-\frac{\kappa\alpha}{4-4\alpha}} M^{-\frac{\beta_\alpha}{2}} + M^{-\frac{1}{2}} + N^{-\frac{1}{2}}.$$

holds, so in order to complete our proof, it suffices to show $\mathbb{E}\|M^{-1}S_M\overline{U}_M\overline{U}_M^T\|$ is of the same order. Note that

$$\overline{U}_M = S_M^{-1}S_{s_1-1}^{e_M}U_{s_1-1} - S_M^{-1}\sum_{i=s_1}^{e_M}\left(S_i^{e_M} + I\right)\eta_i\xi_i$$

so using the fact that ξ_i are martingale differences

$$\begin{aligned} &\mathbb{E}\|\overline{U}_M\|_2^2 \\ \leq\ & S_M^{-2}\|S_{s_1-1}^{e_M}\|^2\mathbb{E}\|U_{s_1-1}\|_2^2 + S_M^{-2}\sum_{i=s_1}^{e_M}\eta_i^2\|S_i^{e_M}+I\|^2\left[\mathrm{tr}(S) + \Sigma_2\mathbb{E}(\|\Delta_i\|_2^\kappa + \|\Delta_i\|_2^2)\right]. \end{aligned}$$

Lemma C.2 claim 3 provides the following bound using the approximations in Lemma C.1

$$M^{-1}S_M^{-1}\|S_{s_1-1}^{e_M}\|^2\mathbb{E}\|U_{s_1-1}\|_2^2 \lesssim M^{-1}S_M^{-1}s_1^\alpha(1+\|U_0\|_2^2 + \|\Delta_0\|_2^2) \lesssim M^{-1}N^{-1}.$$

On the other hand, Lemma C.2 3) also shows that $\eta_i^2\|S_i^{e_M}+I\|^2$ is uniformly bounded, and by Lemma 2.6,

$$[\mathrm{tr}(S) + \Sigma_2\mathbb{E}(\|\Delta_i\|_2^\kappa + \|\Delta_i\|_2^2)] \lesssim 1 + \|\Delta_0\|_2^2$$

by Cauchy-Schwartz inequality. Therefore

$$M^{-1}S_M^{-1}\sum_{i=s_1}^{e_M}\eta_i^2\|S_i^{e_M}+I\|^2\left[\mathrm{tr}(S) + \Sigma_2\mathbb{E}\|\Delta_i\|_2^2\right] \lesssim M^{-1}(1+\|\Delta_0\|_2^2).$$

Collecting the terms with lower order produces the claimed result. ∎

C.4 Consistency proof of Batch-Means Estimator

Proof of Theorem 3.2. Recall the linear oracle sequence defined by (31): $U_n = (I - \eta_n A)\, U_{n-1} + \eta_n \xi_n$, $U_0 = \Delta_0$. From Lemma 3.3

$$\begin{aligned}
&\mathbb{E}\left\| M^{-1} \sum_{k=1}^{M} n_k (\overline{U}_{n_k} - \overline{U}_M)(\overline{U}_{n_k} - \overline{U}_M)^T - A^{-1} S A^{-1} \right\| \\
\lesssim\ & N^{-\frac{\kappa\alpha}{4-4\alpha}} M^{-\frac{\beta_\alpha}{2}} + M^{-\frac{1}{2}} + N^{-\frac{1}{2}}.
\end{aligned}$$

Let $\delta_n := \Delta_n - U_n$, which follows a recursion:

$$\begin{aligned}
\delta_n &= \Delta_{n-1} - U_{n-1} - \eta_n \nabla F(x_{n-1}) - \eta_n A U_{n-1} \\
&= (I - \eta_n A)\delta_{n-1} + \eta_n (A\Delta_{n-1} - \nabla F(x_{n-1})).
\end{aligned}$$

Define the mean differences for batches and overall:

$$\overline{\delta}_{n_k} := n_k^{-1} \sum_{i=s_k}^{e_k} (\Delta_i - U_i), \quad \overline{\delta}_M := S_M^{-1} \sum_{i=s_1}^{e_M} (\Delta_i - U_i).$$

Using $\sum_{k=1}^M n_k \overline{U}_M = \sum_{k=1}^M n_k \overline{U}_{n_k}$ and $\sum_{k=1}^M n_k \overline{X}_M = \sum_{k=1}^M n_k \overline{X}_{n_k}$, we have the following decomposition:

$$\begin{aligned}
&\frac{1}{M}\sum_{k=1}^{M} n_k (\overline{X}_{n_k} - \overline{X}_M)(\overline{X}_{n_k} - \overline{X}_M)^T = \frac{1}{M}\sum_{k=1}^{M} n_k \overline{X}_{n_k}\overline{X}_{n_k}^T - \frac{1}{M}\sum_{k=1}^{M} n_k \overline{X}_M \overline{X}_M^T \\
=&\frac{1}{M}\sum_{k=1}^{M} n_k \overline{\Delta}_{n_k}\overline{\Delta}_{n_k}^T - \frac{1}{M}\sum_{k=1}^{M} n_k x^* x^{*T} + \frac{1}{M}\sum_{k=1}^{M} n_k \overline{\Delta}_M x^{*T} \\
&+ \frac{1}{M}\sum_{k=1}^{M} n_k x^* \overline{\Delta}_M^T - \frac{1}{M}\sum_{k=1}^{M} n_k (x^* + \overline{\Delta}_M)(x^* + \overline{\Delta}_M)^T \\
=&\frac{1}{M}\sum_{k=1}^{M} n_k \overline{\Delta}_{n_k}\overline{\Delta}_{n_k}^T - \frac{1}{M}\sum_{k=1}^{M} n_k \overline{\Delta}_M \overline{\Delta}_M^T \\
=&\frac{1}{M}\sum_{k=1}^{M} n_k (\overline{U}_{n_k} + \overline{\delta}_{n_k})(\overline{U}_{n_k} + \overline{\delta}_{n_k})^T - \frac{1}{M}\sum_{k=1}^{M} n_k (\overline{U}_M + \overline{\delta}_M)(\overline{U}_M + \overline{\delta}_M)^T \\
=&\frac{1}{M}\sum_{k=1}^{M} n_k (\overline{U}_{n_k} - \overline{U}_M)(\overline{U}_{n_k} - \overline{U}_M)^T + \frac{1}{M}\sum_{k=1}^{M} n_k (\overline{U}_{n_k} - \overline{U}_M)(\overline{\delta}_{n_k} - \overline{\delta}_M)^T \\
&+ \frac{1}{M}\sum_{k=1}^{M} n_k (\overline{\delta}_{n_k} - \overline{\delta}_M)(\overline{U}_{n_k} - \overline{U}_M)^T + \frac{1}{M}\sum_{k=1}^{M} n_k (\overline{\delta}_{n_k} - \overline{\delta}_M)(\overline{\delta}_{n_k} - \overline{\delta}_M)^T.
\end{aligned}$$

Then by Young's inequality and $\|C\| \le \mathrm{tr}(C)$ for any PSD matrix C,

$$\mathbb{E}\left\| \frac{1}{M}\sum_{k=1}^{M} n_k (\overline{X}_{n_k} - \overline{X}_M)(\overline{X}_{n_k} - \overline{X}_M)^T - A^{-1} S A^{-1} \right\|$$

$$
\begin{aligned}
\leq & \mathbb{E}\left\|\frac{1}{M}\sum_{k=1}^{M} n_k(\overline{U}_{n_k}-\overline{U}_M)(\overline{U}_{n_k}-\overline{U}_M)^T - A^{-1}SA^{-1}\right\| \\
& + \frac{1}{M}\sum_{k=1}^{M} n_k \mathbb{E}\mathrm{tr}\left[(\overline{\delta}_{n_k}-\overline{\delta}_M)(\overline{\delta}_{n_k}-\overline{\delta}_M)^T\right] \\
& + \frac{2}{M}\sqrt{\sum_{k=1}^{M} n_k \mathbb{E}\mathrm{tr}[(\overline{U}_{n_k}-\overline{U}_M)(\overline{U}_{n_k}-\overline{U}_M)^T]\sum_{k=1}^{M} n_k \mathbb{E}\mathrm{tr}\left[(\overline{\delta}_{n_k}-\overline{\delta}_M)(\overline{\delta}_{n_k}-\overline{\delta}_M)^T\right]}
\end{aligned}
$$

Note that

$$
\begin{aligned}
& \sum_{k=1}^{M} n_k \mathrm{tr}(\overline{\delta}_{n_k}-\overline{\delta}_M)(\overline{\delta}_{n_k}-\overline{\delta}_M)^T \\
= & \sum_{k=1}^{M} n_k\left[\mathrm{tr}(\overline{\delta}_{n_k}\overline{\delta}_{n_k}^T) - \mathrm{tr}(\overline{\delta}_M\overline{\delta}_M^T)\right] \leq \sum_{k=1}^{M} n_k \mathrm{tr}(\overline{\delta}_{n_k}\overline{\delta}_{n_k}^T).
\end{aligned}
$$

Therefore, a proper bound for $\frac{1}{M}\sum_{k=1}^{M} n_k \mathbb{E}\|\overline{\delta}_{n_k}\|_2^2$ will conclude our proof.

Using the definination of Y_j^k and S_j^k in Lemma C.2, we find that

$$
\begin{aligned}
\delta_n &= (I-\eta_n A)\delta_{n-1} + \eta_n(A\Delta_{n-1} - \nabla F(x_{n-1})) \\
&= \left[\prod_{i=s_k}^{n}(I-\eta_i A)\right]\delta_{s_k-1} + \sum_{i=s_k}^{n}\left[\prod_{j=i+1}^{n}(I-\eta_j A)\right]\eta_i(A\Delta_{i-1}-\nabla F(x_{i-1})) \\
&= Y_{s_k-1}^n \delta_{s_k-1} + \sum_{i=s_k}^{n} Y_i^n \eta_i(A\Delta_{i-1}-\nabla F(x_{i-1})) \\
\overline{\delta}_{n_k} &= \frac{1}{n_k}\sum_{j=s_k}^{e_k}\delta_j = \frac{1}{n_k}\sum_{j=s_k}^{e_k} Y_{s_k-1}^j \delta_{s_k-1} + \frac{1}{n_k}\sum_{j=s_k}^{e_k}\sum_{i=s_k}^{j} Y_i^j \eta_i(A\Delta_{i-1}-\nabla F(x_{i-1})) \\
&= \frac{1}{n_k} S_{s_k-1}^{e_k}\delta_{s_k-1} + \frac{1}{n_k}\sum_{i=s_k}^{e_k}\sum_{j=i}^{e_k} Y_i^{e_k}\eta_i(A\Delta_{i-1}-\nabla F(x_{i-1})) \\
&= \frac{1}{n_k} S_{s_k-1}^{e_k}\delta_{s_k-1} + \frac{1}{n_k}\sum_{n=s_k}^{e_k}(I+S_n^{e_k})\eta_n(A\Delta_{n-1}-\nabla F(x_{n-1})).
\end{aligned}
$$

So by Young's inequality and Cauchy-Schwartz inequality

$$
\begin{aligned}
\mathbb{E}\|\overline{\delta}_{n_k}\|_2^2 \leq & 2n_k^{-2}\|S_{s_k-1}^{e_k}\|^2\mathbb{E}\|\delta_{s_k-1}\|_2^2 \\
& + 2n_k^{-2}\mathbb{E}\left(\sum_{n=s_k}^{e_k}\|\eta_n(I+S_n^{e_k})\|\|(A\Delta_{n-1}-\nabla F(x_{n-1})\|\right)^2 \\
\leq & 2n_k^{-2}\|S_{s_k-1}^{e_k}\|^2\mathbb{E}\|\delta_{s_k-1}\|_2^2
\end{aligned}
$$

$$+ 2n_k^{-2} \sum_{n=s_k}^{e_k} \|\eta_n(S_n^{e_k} + I)\|^2 \mathbb{E} \sum_{n=s_k}^{e_k} \|A\Delta_{n-1} - \nabla F(x_{n-1})\|_2^2$$

Apply Lemma 2.6 and Lemma C.2,

$$\|S_{s_k-1}^{e_k}\|^2 \lesssim s_k^{2\alpha}, \quad \sum_{n=s_k}^{e_k} \|\eta_n S_n^{e_k}\|^2 \lesssim n_k.$$

$$\mathbb{E} \sum_{n=s_k}^{e_k} \|A\Delta_{n-1} - \nabla F(x_{n-1})\|_2^2 \lesssim \sum_{n=s_k}^{e_k} \mathbb{E}\|\Delta_{n-1}\|_2^4 \lesssim n_k s_k^{-2\alpha}(1 + \|\Delta_0\|_2^4).$$

Notice that Lemma 2.6 can be applied to both U_n and Δ_n, so

$$\mathbb{E}\|\delta_{s_k-1}\|_2^2 \le 2\mathbb{E}\|\Delta_{s_k-1}\|_2^2 + 2\mathbb{E}\|U_{s_k-1}\|_2^2 \lesssim s_k^{-\alpha}(1 + \|\Delta_0\|_2^2 + \|U_0\|_2^2).$$

Furthermore, recall the batch size $e_k = ((k+1)N)^{\frac{1}{1-\alpha}}$,

$$n_k = \left[(k+1)^{\frac{1}{1-\alpha}} - k^{\frac{1}{1-\alpha}}\right] N^{\frac{1}{1-\alpha}} = \frac{N^{\frac{1}{1-\alpha}}}{1-\alpha} \int_k^{k+1} x^{\frac{\alpha}{1-\alpha}} dx \asymp k^{\frac{\alpha}{1-\alpha}} N^{\frac{1}{1-\alpha}},$$

$$n_k \mathbb{E}\|\overline{\delta}_{n_k}\|_2^2 \lesssim n_k^{-1} s_k^{\alpha} + n_k s_k^{-2\alpha} \lesssim N^{-1} + k^{-\frac{\alpha}{1-\alpha}} N^{\frac{1-2\alpha}{1-\alpha}}.$$

For the sum of geometric sequence in the last term, the following holds

$$\frac{1}{M} \sum_{k=1}^{M} k^{-\frac{\alpha}{1-\alpha}} \lesssim \frac{1}{M} \int_1^{M+1} x^{-\frac{\alpha}{1-\alpha}} dx \lesssim M^{(1-\frac{\alpha}{1-\alpha})^+ - 1},$$

which implies

$$\frac{1}{M} \sum_{k=1}^{M} n_k \mathbb{E}\|\overline{\delta}_{n_k}\|_2^2 \lesssim N^{-1} + M^{(\frac{1-2\alpha}{1-\alpha})^+ - 1} N^{\frac{1-2\alpha}{1-\alpha}}.$$

In conclusion,

$$\begin{aligned} &\mathbb{E}\left\| M^{-1} \sum_{k=1}^{M} n_k (\overline{X}_{n_k} - \overline{X}_M)(\overline{X}_{n_k} - \overline{X}_M)^T - A^{-1} S A^{-1} \right\| \\ \lesssim & N^{-\frac{\kappa\alpha}{4-4\alpha}} M^{-\frac{\beta_\alpha}{2}} + M^{-\frac{1}{2}} + N^{-\frac{1}{2}} + M^{(\frac{1-2\alpha}{2-2\alpha})^+ - \frac{1}{2}} N^{\frac{1-2\alpha}{2-2\alpha}}. \end{aligned}$$

□

D Proofs for results in the High Dimensional Section (Section 4)

D.1 Proof of Proposition 4.1

Proof. The convergence of $\|\widehat{x}_n - x^*\|_1$ in (42) is a direct consequence of Proposition 1 in Agarwal et al. [2012]. To see this, from Lemma 1 and Proposition 1 of Agarwal et al.

[2012], By Lemma 1 (Eq. (39)) of Agarwal et al. [2012], we have

$$\|\widehat{x}_n - x^*\|_1 \lesssim s\lambda_{K_n}, \tag{56}$$

where K_n is the number of epochs and λ_{K_n} is the regularization parameters used in the K_n-th epoch (see Eq. (34)). By the setting of λ_i for the i-th epoch in Eq. (34) in Agarwal et al. [2012], we have

$$\lambda_i^2 = c\frac{R_i}{s\sqrt{T_i}}\sqrt{\log d}, \tag{57}$$

where $R_i = \frac{R_1}{2^{i/2}}$ and the number of iterations in the i-th epoch T_i is set to

$$T_i \geq c_1 \frac{s^2}{R_i^2}\log d, \tag{58}$$

according to Eq. (32) in Agarwal et al. [2012]. Combine (57) and (58), we have

$$\lambda_{K_n} \lesssim \frac{R_{K_n}}{s} \lesssim \frac{1}{2^{K_n/2}s} \tag{59}$$

To further bound λ_{K_n} in (59), we only need an lower bound on the number of epochs K_n. By the setting of T_i in Eq. (32) in Agarwal et al. [2012], we can easily derive that,

$$n = \sum_{i=1}^{K_n} T_i \lesssim s^2(\log d)2^{K_n},$$

which implies that

$$\lambda_{K_n} \lesssim \frac{1}{2^{K_n/2}s} \lesssim \sqrt{\frac{\log d}{n}},$$

which combining with (56) completes the proof of (42).

For (43), recall that $\widehat{\gamma}^j$ is the solution of RADAR for solving the problem in (39). We can apply the Proposition 1 from [Agarwal et al., 2012] to provide the bound on $\|\widehat{\gamma}^j - (\gamma^j)^*\|_1$. Here $(\gamma^j)^*$ is the optimum for the un-regularized population loss, i.e.,

$$(\gamma^j)^* = \operatorname{argmin} \mathbb{E}\left((a_j - a_{-j}\gamma^j)^2\right) \tag{60}$$

According to the normal assumption of the design $a \sim N(0, A)$, the conditional distribution of a_j given a_{-j} is

$$a_j | a_{-j} \sim N\left(-\Omega_{j,j}^{-1}\Omega_{j,-j}a_{-j}, \Omega_{j,j}^{-1}\right),$$

which implies that

$$(\gamma^j)^* = -\Omega_{j,j}^{-1}\left(\Omega_{j,-j}\right)^T.$$

We further note that by the row sparsity assumption in (36), $(\gamma^j)^*$ is an s_j-sparse vector. By exactly the same argument as for proving (56), we obtain the result in (43). □

D.2 Proof of Theorem 4.2

We first prove two lemmas about estimating Ω_{jj} and Ω,

Lemma D.1. *Let* $\hat{\tau}_j \equiv \frac{1}{n}\|D_j - D_{-j}\hat{\gamma}_j\|_2^2$. *Assuming* $\|\hat{\gamma}_j - \gamma_j\|_2 \le C\sqrt{\frac{s_j \log d}{n}}$, $\|\hat{\gamma}_j - \gamma_j\|_1 \le Cs_j\sqrt{\frac{\log d}{n}}$, *and* $n > s_j \log d$, *we have*

$$|\frac{\hat{\tau}_j}{\Omega_{jj}^{-1}} - 1| \le C\sqrt{\frac{\log d}{n}} + 2\lambda_{\max}(A)C\frac{s_j \log d}{n} \tag{61}$$

Proof. Recall that $a \sim \mathcal{N}(0, A)$, so $a_j = -\Omega_{jj}^{-1}\Omega_{j,-j}a_{-j} + \sqrt{\Omega_{jj}^{-1}}z_j = \gamma_j^T a_{-j} + \sqrt{\Omega_{jj}^{-1}}z_j$ where $z_j \sim \mathcal{N}(0,1)$.

$$\begin{aligned}\sqrt{\hat{\tau}_j} = \frac{1}{\sqrt{n}}\|D_j - D_{-j}\hat{\gamma}_j\|_2 &= \|D_j - D_{-j}\gamma_j + D_{-j}(\hat{\gamma}_j - \gamma_j)\|_2 \\ &= \frac{1}{\sqrt{n}}\left\|\sqrt{\Omega_{jj}^{-1}}z_j + D_{-j}(\hat{\gamma}_j - \gamma_j)\right\|_2 \\ &\le \frac{1}{\sqrt{n}}\left\|\sqrt{\Omega_{jj}^{-1}}z_j\right\|_2 + \frac{1}{\sqrt{n}}\|D_{-j}(\hat{\gamma}_j - \gamma_j)\|_2 \end{aligned} \tag{62}$$

Using Loh et al. [2012, Equation 2.13] , we have that

$$\begin{aligned}\frac{1}{n}\|D_{-j}(\hat{\gamma}_j - \gamma_j)\|^2 &\le 2\lambda_{\max}(A)\|\hat{\gamma}_j - \gamma_j\|_2^2 + C\frac{\log d}{n}\|\hat{\gamma}_j - \gamma_j\|_1^2 \\ \frac{1}{\sqrt{n}}\|D_{-j}(\hat{\gamma}_j - \gamma_j)\| &\le \sqrt{2\lambda_{\max}(A)}\|\hat{\gamma}_j - \gamma_j\|_2 + \sqrt{C\frac{\log d}{n}}\|\hat{\gamma}_j - \gamma_j\|_1\end{aligned}$$

By the concentration inequality for χ^2 , $\Pr(\|z_j\|^2 > n + 2\sqrt{nt} + 2t) \le e^{-t}$, so letting $t \asymp C\log d$,

$$\begin{aligned}\Pr(\|z_j\|^2 > n + 2\sqrt{n\log d} + 2\log d) &\le \frac{1}{d^C} \\ \Pr(\frac{\|z_j\|}{\sqrt{n}} - 1 > \sqrt{2\sqrt{\log d/n} + 2\log d/n}) &\le \frac{1}{d^C}\end{aligned}$$

Using these two estimates with (62) and $\frac{\log d}{n} < 1$,

$$\begin{aligned}\sqrt{\hat{\tau}_j} &\le \sqrt{\Omega_{jj}^{-1}}\left(1 + \sqrt{2\sqrt{\log d/n} + 2\log d/n}\right) \\ &\quad + \sqrt{2\lambda_{\max}(A)}\|\hat{\gamma}_j - \gamma_j\|_2 + \sqrt{C\frac{\log d}{n}}\|\hat{\gamma}_j - \gamma_j\|_1 \\ &\le \sqrt{\Omega_{jj}^{-1}}(1 + C\left(\frac{\log d}{n}\right)^{.25}) + \sqrt{2\lambda_{\max}(A)}\sqrt{\frac{s_j\log d}{n}} + Cs_j\frac{\log d}{n}.\end{aligned}$$

By a similar argument to (62) and using concentration inequality for χ^2, $\Pr(\|z_j\|^2 < n - 2\sqrt{nt}) \le e^{-t}$,

$$\begin{aligned}\sqrt{\hat{\tau}_j} &\ge \frac{1}{\sqrt{n}}\left\|\sqrt{\Omega_{jj}^{-1}}z_j\right\|_2 - \frac{1}{\sqrt{n}}\|D_{-j}(\hat{\gamma}_j - \gamma_j)\|_2 \\ &\ge \sqrt{\Omega_{jj}^{-1}}(1 - C\left(\frac{\log d}{n}\right)^{.25}) - \sqrt{2\lambda_{\max}(A)}\sqrt{\frac{s_j \log d}{n}} - Cs_j\frac{\log d}{n}\end{aligned}$$

Dividing by $\sqrt{\Omega_{jj}^{-1}}$, squaring, and using $\frac{s_j \log d}{n} < 1$ gives

$$|\frac{\hat{\tau}_j}{\Omega_{jj}^{-1}} - 1| \lesssim \sqrt{\frac{\log d}{n}} + 2\lambda_{\max}(A)\frac{s_j \log d}{n}.$$

Lemma D.2. *Under the Gaussian assumption of the design, the row sparsity assumption of Ω in (36) and the assumption on s_j in Theorem 4.2, we have with high probability*

$$\max_{j\in\{1,\ldots,d\}} \|\widehat{\Omega}_j - \Omega_j\|_2 \le C\sqrt{\frac{\max_j s_j \log d}{n}}$$
$$\max_{j\in\{1,\ldots,d\}} \|\widehat{\Omega}_j - \Omega_j\|_1 \le C \max_j s_j \sqrt{\frac{\log d}{n}}$$

□

Proof of Lemma D.2. By the definition of $\widehat{\Omega}$ in (41), we have the j-th row of $\widehat{\Omega}$ takes the following form,

$$\widehat{\Omega}_j = \widehat{C}_j/\widehat{\tau}_j, \qquad \widehat{C}_j = (-\widehat{\gamma}_1^j, \ldots, -\widehat{\gamma}_{j-1}^j, 1, \ldots, -\widehat{\gamma}_d^j).$$

By Proposition 4.1, we have for any j,

$$\|\widehat{C}_j - \Omega_{j,j}^{-1}\Omega_j\|_1 \lesssim s_j\sqrt{\frac{\log d}{n}}, \tag{63}$$

holding with high probability.

Using Lemma D.1, $\hat{\tau}_j = O(1)$, and noting that $\frac{s_j \log d}{n} < \sqrt{\frac{\log d}{n}}$, we see that

$$\left|\frac{\Omega_{jj}^{-1}}{\hat{\tau}_j} - 1\right| \le C\sqrt{\frac{\log d}{n}}. \tag{64}$$

Combining (64) with (63) and and noticing that

$$\|\Omega_j\|_1 \le \sqrt{s_j}\|\Omega_j\|_2 \le \frac{\sqrt{s_j}}{\lambda_{\min}(A)},$$

we have with probability at least $1 - 1/d^2$,

$$\begin{aligned}\|\widehat{\Omega}_j - \Omega_j\|_1 &= \left\| \left(\widehat{C}_j/\widehat{\tau}_j - \Omega_{j,j}^{-1}\Omega_j/\widehat{\tau}_j \right) + \left(\Omega_{j,j}^{-1}\Omega_j/\widehat{\tau}_j - \Omega_j \right) \right\|_1 \\ &\leq \frac{1}{\widehat{\tau}_j}\|\widehat{C}_j - \Omega_{j,j}^{-1}\Omega_j\|_1 + \left| \frac{\Omega_{j,j}^{-1}}{\widehat{\tau}_j} - 1 \right| \|\Omega_j\|_1 \\ &\lesssim s_j \sqrt{\frac{\log d}{n}},\end{aligned}$$

which completes the proof by taking a union bound over j. The proof for the l_2 bound is analogous. □

With Lemma D.2 in place, we are ready to prove Theorem 4.2.

Proof. We first plug the linear model $b = Dx^* + \epsilon$ into the definition of the the desparsified estimator in (38) and obtain,

$$\begin{aligned}\widehat{x}^d &= \widehat{x}_n + \frac{1}{n}\widehat{\Omega}\widehat{A}\left(x^* - \widehat{x}_n\right) + \frac{1}{n}\widehat{\Omega}D^T\epsilon \\ &= x^* + \left(\widehat{\Omega}\widehat{A} - I\right)\left(x^* - \widehat{x}_n\right) + \frac{1}{n}\widehat{\Omega}D^T\epsilon.\end{aligned} \tag{65}$$

It is easy to see that $\frac{1}{\sqrt{n}}\widehat{\Omega}D^T\epsilon \sim N(0, \sigma^2\widehat{\Omega}\widehat{A}\widehat{\Omega}^T)$. Therefore, to prove Theorem 4.2, we only need to show the term

$$\Delta := \left(\widehat{\Omega}\widehat{A} - I\right)\left(x^* - \widehat{x}_n\right) = o_p(\frac{1}{\sqrt{n}}), \tag{66}$$

which gives the result in (44) by Slutsky's theorem.

By adding and subtracting $\Omega\hat{A}(\hat{x}_n - x^*)$,

$$\begin{aligned}\|\Delta\|_\infty &\leq \left\|(\widehat{\Omega}\hat{A} - \Omega\hat{A})(\hat{x}_n - x^*)\right\|_\infty + \left\|(\Omega\hat{A} - I)(\hat{x}_n - x^*)\right\|_\infty \\ &\leq \left\|(\widehat{\Omega}\hat{A} - \Omega\hat{A})(\hat{x}_n - x^*)\right\|_\infty + \left\|\Omega\hat{A} - I\right\|_\infty \|\hat{x}_n - x^*\|_1 \\ &\leq \left\|(\widehat{\Omega}\hat{A} - \Omega\hat{A})(\hat{x}_n - x^*)\right\|_\infty + C\frac{s\log d}{n}\end{aligned}$$

where in the last inequality we used [Javanmard and Montanari, 2014, Lemma 23] and Proposition 4.1.

The first term can be bounded

$$
\begin{aligned}
&\left\|(\widehat{\Omega}\hat{A}-\Omega\hat{A})(\hat{x}_n-x^*)\right\|_\infty \\
=&\left\|(\widehat{\Omega}-\Omega)\frac{1}{n}D^TD(\hat{x}_n-x^*)\right\|_\infty \\
\leq&\left[\max_j \frac{1}{\sqrt{n}}\left\|D(\widehat{\Omega}^T-\Omega)e_j\right\|_2\right]\frac{1}{\sqrt{n}}\left\|D(\hat{x}_n-x^*)\right\|_2 \\
\leq&\left[\max_j \sqrt{C_1\lambda_{\max}(A)\left\|(\widehat{\Omega}^T-\Omega)e_j\right\|_2^2+C_2\frac{\log d}{n}\left\|(\widehat{\Omega}^T-\Omega)e_j\right\|_1^2}\right] \\
&\sqrt{C_1'\lambda_{\max}(A)\left\|\hat{x}_n-x^*\right\|_2^2+C_2'\frac{\log d}{n}\left\|\hat{x}_n-x^*\right\|_1^2} \\
\lesssim&\max_j\sqrt{\frac{s_j\log d}{n}+\frac{s_j^2\log^2 d}{n^2}}\sqrt{\frac{s\log d}{n}+\frac{s^2\log^2 d}{n^2}} \\
\lesssim&s\frac{\log d}{n}
\end{aligned}
$$

where in the second inequality we used [Loh et al., 2012, Equation 2.13], and in the third inequality we used

$$
\left\|(\widehat{\Omega}^T-\Omega)e_j\right\|_2\lesssim\sqrt{\frac{s_j\log d}{n}},\quad \left\|(\widehat{\Omega}^T-\Omega)e_j\right\|_1\lesssim Cs_j\sqrt{\frac{\log d}{n}}
$$

from Lemma D.2, and $\|\hat{x}_n-x^*\|_2\lesssim C\sqrt{\frac{s\log d}{n}}$, and $\|\hat{x}_n-x^*\|_1\lesssim s\sqrt{\frac{\log d}{n}}$ from Agarwal et al. [2012] and Proposition 4.1. In the last inequality, we use $s=o(\sqrt{n}/\log d)$ and $s_j\leq Cs$.

By the assumption that $s=o(\sqrt{n}/\log d)$, we have that $\|\Delta\|_\infty\leq o(\frac{1}{\sqrt{n}})$, which completes the proof.

□